

*Final
Report
2007*

Final Report

Decision-Aiding and Optimization for Vertical Navigation of Long-Haul Aircraft

Nicholas J.M. Patrick
Thomas B. Sheridan

Human-Machine Systems Laboratory
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

August 1996

Submitted to the National Aeronautics and Space
Administration, under contract number NAG 2-729.

Decision-Aiding and Optimization for Vertical Navigation of Long-Haul Aircraft

by

Nicholas J.M. Patrick

Submitted to the Department of Mechanical Engineering on August 6, 1996
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Mechanical Engineering

Abstract

Most decisions made in the cockpit are related to safety, and have therefore been proceduralized in order to reduce risk. There are very few which are made on the basis of a value metric such as economic cost. One which can be shown to be value based, however, is the selection of a flight profile.

Fuel consumption and flight time both have a substantial effect on aircraft operating cost, but they cannot be minimized simultaneously. In addition, winds, turbulence, and performance vary widely with altitude and time. These factors make it important and difficult for pilots to (a) evaluate the outcomes associated with a particular trajectory before it is flown and (b) decide among possible trajectories. The two elements of this problem considered here are (i) determining what constitutes optimality, and (ii) finding optimal trajectories.

Pilots and dispatchers from major U.S. airlines were surveyed to determine which attributes of the outcome of a flight they considered the most important. Avoiding turbulence—for passenger comfort—topped the list of items which were not safety related. Pilots' decision making about the selection of flight profile on the basis of flight time, fuel burn, and exposure to turbulence was then observed. Of the several behavioral and prescriptive decision models invoked to explain the pilots' choices, utility maximization is shown to best reproduce the pilots' decisions.

After considering more traditional methods for optimizing trajectories, a novel method is developed using a genetic algorithm (GA) operating on a discrete representation of the trajectory search space. The representation is a sequence of command altitudes, and was chosen to be compatible with the constraints imposed by Air Traffic Control, and with the training given to pilots. Since trajectory evaluation for the GA is performed holistically, a wide class of objective functions can be optimized easily. Also, using the GA it is possible to compare the costs associated with different airspace design and air traffic management policies.

A decision aid is proposed which would combine the pilot's notion of optimality with the GA-based optimization, provide the pilot with a number of alternative pareto-optimal trajectories, and allow him to consider unmodelled attributes and constraints in choosing among them. A solution to the problem of displaying alternatives in a multi-attribute decision space is also presented.

Thesis Supervisor: Thomas B. Sheridan
Title: Ford Professor of Engineering and Applied Psychology
and Professor of Aeronautics and Astronautics

List of Abbreviations

ACARS	Airborne Communication and Reporting System
AGL	Above Ground Level
AI	Artificial Intelligence
AIM	Aeronautical Information Manual
AOC	Airline Operational Control
ATC	Air Traffic Control
ATP	Airline Transport Pilot
DM	Decision Maker
EMV	Expected Monetary Value
EU	Expected Utility
EV	Expected Value
EVPI	Expected Value of Perfect Information
FAA	Federal Aviation Administration
FAR	Federal Aviation Regulations
FIR	Flight Information Region
FL	Flight Level
FMC	Flight Management Computer
FMS	Flight Management System
GA	Genetic Algorithm
GPS	Global Positioning System
GPWS	Ground Proximity Warning System
HF	High Frequency
ICAO	International Civil Aviation Organization
IFR	Instrument Flight Rules
IMC	Instrument Meteorological Conditions
kft	Thousands of feet
klb	Thousands of pounds
kt	A knot, which is one nautical mile per hour
LRC	Long Range Cruise
MAU	Multi-Attribute Utility
MIDAS	Man-Machine Integration Design and Analysis System
MSL	Above Mean Sea Level
NASA	National Aeronautics and Space Administration
PDF	Probability Density Function
PIREP	Pilot Report
SATCOM	Satellite Communications
SEU	Subjectively Expected Utility
SIGMET	A safety-related weather advisory
TOGW	Take-Off Gross Weight
UI	Utility Independent
VHF	Very High Frequency

Notation

Decision Theory

$A \succ B$	DM prefers outcome A to outcome B
$A \prec B$	DM prefers outcome B to outcome A
$A \sim B$	DM is indifferent to outcomes A and B
$A \preceq B$	DM does not prefer outcome A to outcome B
\mathcal{O}	An objective function to be maximized or minimized by the GA
\mathbf{R}^m	An m -dimensional space
$r_x(x_a)$	Risk aversion over attribute x , at $x = x_a$
x, y, z	Scalar attributes
x_0	Value of x such that $\mathcal{U}_x(x_0) = 0$
x_1	Value of x such that $\mathcal{U}_x(x_1) = 1$
\hat{x}	An estimate of x
\mathbf{x}	A vector of attributes: $\mathbf{x} \in \mathbf{R}^m$
$\mathcal{U}_x(x_a)$	Conditional utility at $x = x_a$
$\mathcal{U}_{xyz}(x_a, y_a, z_a)$	Utility at $x = x_a$, $y = y_a$, and $z = z_a$
\mathcal{U}_{101}	Utility “corner” point with, for example, $x = x_1$, $y = y_0$, $z = z_1$
$\mathcal{V}_x(x_a)$	Value function over x at $x = x_a$
Ω	The set of feasible points in the solution space: $\Omega \subset \mathbf{R}^m$
$\langle x_a, p; x_b \rangle$	A lottery whose outcome is x_a with probability p , or x_b with probability $1 - p$

Statistics

\bar{R}	Mean rank
t	Student’s t statistic
W	Kendal coefficient of concordance
ν	Number of degrees of freedom

Navigation

A	Altitude (in thousands of feet)
a	Wind correction angle (in degrees)
c	True bearing of wind (in degrees clockwise from True North)
d	Great-circle distance between two waypoints (in nautical miles)
EHW	Effective headwind (in knots)
GS	Ground speed (in knots)
h	True heading (in degrees clockwise from True North)
$\mathbf{i}, \mathbf{j}, \mathbf{k}$	Unit vectors in an Earth-centered cartesian coordinate system
R_e	Radius of the Earth (in nautical miles)
TAS	True airspeed (in knots)
\mathbf{u}	A unit vector
γ	Angle subtended at the Earth's center by two waypoints (in degrees)
ϕ	Longitude (in degrees east of the Prime Meridian)
ψ	Latitude (in degrees north of the Equator)
ω	True course (in degrees clockwise from True North)
Δ	A small change in a parameter or variable

Aircraft Performance

C, C'	Cost (in dollars and thousands of pounds of fuel, respectively)
c_f	Cost of fuel (in cents per pound)
c_t	Aircraft operating cost (excluding fuel, in dollars per hour)
D	Distance flown (in nautical miles)
f	Fuel burn (in thousands of pounds)
I_c	Cost index
m	Aircraft mass (in thousands of pounds)
n_c	Number of climbs in a flight
n_d	Number of descents in a flight
q	A measure of ride quality or passenger comfort
t	Flight time (in hours)
Δt	Integration interval (in hours)

Contents

1	Introduction	19
1.1	Motivation	19
1.2	Objectives	19
1.3	The Profile Selection Problem	20
1.3.1	Decision Makers	20
1.3.2	Search Space	20
1.4	Outline	22
2	Decision Models: Prescriptive and Behavioral	25
2.1	A Framework for Decision Modelling	25
2.2	Simple Decision Models	26
2.2.1	Behavioral Models	26
2.2.2	Proceduralized Decision Making	27
2.2.3	Prescriptive Models	27
2.3	Utility Theory	28
2.3.1	Axiomatic Basis	28
2.3.2	Model Decisions About Performance, Not Safety	28
2.3.3	Prescriptive Use of Utility Theory	29
3	Pilot and Dispatcher Surveys	31
3.1	An Exploratory Survey of Pilots	31
3.2	A Forced-Response Survey of Pilots and Dispatchers	32
3.2.1	Pilots' Rankings of Attributes of the Outcome	32
3.2.2	Dispatchers' Rankings of Attributes of the Outcome	34
3.2.3	Agreement Between Pilots and Dispatchers	34
3.2.4	Pilots' Rankings of Sources of Information	35
3.3	FMS Features Desired by Pilots	36
4	Utility Assessment Experiment: Eliciting Pilots' Objective Functions	37
4.1	Background	37
4.2	Verification and Scaling of Important Attributes	37
4.3	Comparisons of Pairs of Alternatives	38
4.4	Using a Multilinear Utility Function	38
4.5	Evaluation of the Conditional Utilities	40
4.5.1	The Importance of Resource Constraints	41
4.6	Fitting the Conditional Utility Functions	42
4.6.1	Linear Utility Function	43
4.6.2	Piecewise Linear, or Interpolated Utility Function	43
4.6.3	Least-Squares Exponential Utility Function	43

4.6.4	Choosing Between Conditional Utility Models	44
4.7	Verification of the Independence of Conditional Utilities	44
4.8	Determination of the Scaling Constants, $k_x \dots k_{xyz}$	45
4.8.1	Graphical Method	45
4.8.2	Lottery Method	45
4.8.3	Calculation of the Scaling Constants	46
4.9	The Complete Multi-Attribute Utility Function	46
4.10	Risk Aversion	47
4.11	A Comparison of Several Decision Models	50
5	Trajectory Optimization: Maximizing the Objective Function	53
5.1	A Survey of Methods	53
5.1.1	Gradient-Based Techniques	53
5.1.2	Filtered Cruise-Climb Solutions	54
5.1.3	A Discrete Representation	54
5.1.4	Optimization by Search	54
5.1.5	Stochastic Directed Search and the Genetic Algorithm	56
5.2	Optimization by Genetic Algorithm	58
5.3	Results	61
5.3.1	Excessive Climbing and Descending	61
5.3.2	Schedule Adherence	62
5.3.3	Airspace Design	63
6	A Decision Support System for Trajectory Optimization	69
6.1	The Current Situation	69
6.2	Design Considerations	69
6.3	The Parallel-Axis Display	70
7	Conclusions, Contributions, and Suggestions for Further Research	75
7.1	Conclusions & Contributions	75
7.1.1	Decision Theory	75
7.1.2	Optimization	76
7.1.3	Airspace and Operations	76
7.2	Suggestions for Further Research	76
7.2.1	Decision Theory	76
7.2.2	Decision Aiding	76
7.2.3	Optimization	76
A	Preliminary Decision Experiment	79
A.1	Background	79
A.2	Method	79
A.2.1	Paired Comparisons	79
A.2.2	Utility Assessment	80
A.2.3	Discriminant-Function Models	82
A.3	Results	83
A.3.1	Error Rates	83
A.3.2	Subjective Difference	83
A.4	Conclusions	84

B	Survey Data	85
B.1	Exploratory Survey	85
B.1.1	FMS Features Desired by Pilots	86
B.2	Pilot-Dispatcher Survey	89
C	Utility Assessment Experiment Data	91
C.1	Conditional Utilities	92
C.2	Multilinear Utilities	97
C.2.1	Multilinear Utility Functions	97
C.2.2	Graphical Comparison of Utility and Cost Models	97
D	Navigation Algorithms	101
D.1	Coordinate Systems	101
D.2	Great-Circle Distance	102
D.2.1	Spherical-Earth Calculation	102
D.2.2	Ellipsoidal Error	102
D.2.3	Altitude-related Error	102
D.3	Great-Circle Course	103
D.4	Heading and groundspeed in the presence of wind	104
D.4.1	True heading	104
D.4.2	Groundspeed	105
D.4.3	Effective headwind	105
E	Aircraft Performance Simulation	107
E.1	Altitude Capability	107
E.2	Climb	107
E.3	Cruise	108
E.4	Descent	109
E.5	Implementation	109
E.5.1	Programming	109
E.5.2	Choice of Integration Interval	110
	References	117

List of Figures

1-1	The engineer's hierarchy of research objectives, in which the engineering problem drives both the scientific questions and the experimental objectives. Note that a physicist might reverse the precedence of the engineering and science items.	20
1-2	Los Angeles to Sydney: a typical trans-Pacific route.	21
1-3	The decision-making trinity in air carrier operations.	21
1-4	The present vertical structure of U.S. Airspace. Note that below 29 kft opposing traffic is separated by one thousand feet, whereas above 29 kft it is separated by two thousand feet. Class A (IFR-only) airspace extends from 18,000 feet MSL to 60,000 feet MSL.	22
1-5	Fuel economy (in nautical air miles per klb of fuel) against aircraft weight and altitude showing the large variation in cruise performance over the aircraft's operating envelope. Data are at Long-Range Cruise (LRC) for the Boeing 747-400 with PW4056 engines.	23
2-1	A decision making paradigm with three operations: <i>selecting</i> admissible solutions (represented here by gray dots) in the <i>search space</i> ($s_1 \dots s_3$), <i>evaluating</i> each candidate to determine its attribute location (x, y) in the <i>decision space</i> , and finally <i>comparing</i> the candidate solutions, by mapping them all onto a single <i>value</i> or <i>utility scale</i> (v).	26
2-2	Hypothetical examples of two types of utility, $U(x)$, as functions of outcome, x , (a) for a safety-related task, and (b) for a performance-related task. Note the avoidance states at one extreme of x in (a).	29
3-1	The attributes of the outcome mentioned in the first (free-response) questionnaire by the 32 pilots, and the fraction of pilots who mentioned each one.	32
3-2	A quasi-cardinal representation of the mean ranks assigned by both pilots and dispatchers to each of the six attributes of the outcome.	35
4-1	The choice between the lottery $\langle x_a, 50\%; x_b \rangle$ and the certain outcome x_c offered to the subjects.	40
4-2	One subject's conditional utilities for each of the three attributes, showing the exponential fit of Equation 4.7. (S_4)	41
4-3	One subject's conditional utilities for time, both with and without a schedule constraint at 15 hours. (S_1)	42
4-4	Three models for a conditional utility function: (a) linear, (b) piecewise linear interpolation, and (c) exponential.	42

4-5	A typical subject's conditional utility for time (S_3), showing both methods of fitting a utility function to the elicited points: (a) linear interpolation, and (b) an exponential function. A straight-line (c), which represents a hypothetical <i>risk-neutral</i> utility function, is shown for comparison.	44
4-6	Graphical assessment of the corner utility points. The subjects were shown the scale in the left hand panel, and were asked to fill it out, producing something like the scale shown in the right hand panel.	45
4-7	Three views of the final multilinear utility function for a typical subject (S_5). For each plot, the omitted variable is held constant at its most desirable value.	47
4-8	Iso-preference curves for two different decision models over flight time and destination fuel for a typical subject (S_5): the multilinear utility function (shown in two dimensions, with turbulence held constant at zero) on the left, and the cost-minimization model on the right. The iso-preference curves are marked in <i>utils</i> and thousands of dollars respectively.	48
4-9	Risk aversions by subject and attribute. Note the risk-seeking behavior of some of the pilots in their preference structure for exposure to turbulence.	49
4-10	A comparison of the predictive performance of the three decision models: utility maximization, lexicographic ordering, and minimization of linear cost. The horizontal gray lines mark the sample means, the vertical gray lines mark the extent of the 95% confidence intervals for the population means; the black boxes mark the 25 th , 50 th , and 75 th percentile points for the samples; and the whiskers mark the 0 th and 100 th percentile points for the samples. In some cases, several sample points were at the bottom of the range, so that the 0 th and 25 th percentile points were coincident, hiding some of the whiskers.	51
5-1	An illustration of an optimal cruise-climb trajectory (labeled <i>raw</i>), and the result of filtering (labeled <i>filtered</i>) by using the <i>nearest</i> allowable eastbound flight level. . . .	54
5-2	Altitude capability as a function of weight for the Boeing 747-400 with P&W 4056 engines.	57
5-3	A flow chart for the simple genetic algorithm.	58
5-4	The chromosome used to represent a typical profile for the fifteen-segment flight. Note that the sequence of command altitudes is read from left to right.	59
5-5	Single-point crossover of two chromosomes, A and B , to produce two offspring, A' and B'	60
5-6	Point mutation of a chromosome, C , to produce a new chromosome, C'	60
5-7	Block mutation of a chromosome, D , to produce a new chromosome, D'	60
5-8	The minimum-fuel ($\mathcal{O} = f$) trajectory produced by the GA. Note the winds—shown in gray—each acting over the third of the route up to its vertical gray line. Note also the difference between the commanded profile in gray, and the profile actually flown in black.	62
5-9	The minimum-time ($\mathcal{O} = t$) trajectory produced by the GA.	63
5-10	The minimum-cost trajectory ($\mathcal{O} = f + \frac{I_c}{10} t$, $I_c = 100$) produced by the GA.	64
5-11	The minimum-cost trajectory ($\mathcal{O} = f + \frac{I_c}{10} t$, $I_c = 100$) with no winds produced by the GA.	65
5-12	The maximum-utility (for S_5) trajectory produced by the GA.	65
5-13	The minimum-cost trajectory with limited climbs and descents ($\mathcal{O} = f + \frac{I_c}{10} t + 0.5(n_c + n_d - 2)$, $I_c = 100$) produced by the GA.	66
5-14	The minimum fuel trajectory with schedule adherence ($\mathcal{O} = f + 30 \text{ abs}[t - 16]$) produced by the GA.	66

5-15	The minimum-cost ($\mathcal{O} = f + \frac{I_c}{10} t$, $I_c = 100$.) trajectory using 2000 foot separations.	67
5-16	The minimum-cost ($\mathcal{O} = f + \frac{I_c}{10} t$, $I_c = 100$.) trajectory using ICAO separations. . . .	67
6-1	A two-dimensional plot of the three-attribute options available to the decision maker.	71
6-2	A parallel-axis plot, which is capable of showing the levels of attainment of many attributes.	72
6-3	A second prototypic parallel-axis plot, in which the axes have been sized according to the relative importance of each attribute.	72
6-4	A third prototypic parallel-axis utility plot, in which the axes have been sized according to the relative importance of each attribute, and distorted to show the nonlinearity of the subject's preference over each attribute.	73
7-1	The augmented chromosome proposed to represent both altitude and speed for a profile. Note that the speed and altitude parts of the chromosome would be manipulated separately—they are only shown together for clarity.	77
A-1	A graphical depiction of a typical situation presented to the subjects. Note that the alternates were always located on the dashed circle, so as to be equally inconvenient for subsequent travel to the ultimate destination.	80
A-2	The scales used in the utility assessment, as returned by a typical subject.	81
A-3	A perceptron, or single-element neural network, which was trained to find a discriminant function (hyperplane) to fit each subject's paired-comparison data.	82
C-1	Conditional utilities for subject one (B-747-400).	92
C-2	Conditional utilities for subject two (MD-11).	93
C-3	Conditional utilities for subject three (B-767).	94
C-4	Conditional utilities for subject four (B-747-400).	95
C-5	Conditional utilities for subject five (B-747-400).	96
C-6	A comparison of utility (with turbulence held constant at zero) and cost models for subject one (B-747-400).	97
C-7	A comparison of utility (with turbulence held constant at zero) and cost models for subject two (MD-11).	98
C-8	A comparison of utility (with turbulence held constant at zero) and cost models for subject three (B-767).	98
C-9	A comparison of utility (with turbulence held constant at zero) and cost models for subject four (B-747-400).	99
C-10	A comparison of utility (with turbulence held constant at zero) and cost models for subject five (B-747-400).	99
D-1	The Earth-centered coordinate systems, showing great-circle distance d , and initial true course ω , from waypoint 1 to waypoint 2.	101
D-2	Altitude-related error.	103
D-3	The wind triangle, showing the relationship between airspeed, windspeed, and ground-speed. Note that all angles are measured <i>clockwise</i> from <i>True</i> North, and that the wind angle, c , is the direction the wind comes <i>from</i>	104
E-1	Altitude capability as a function of weight. The solid and the continuous lines represent cubic least-squares lines fitted to the data from the performance manual for two different temperatures, while the stepped function represents the maximum altitude used in the simulation.	108

E-2	Time to climb and adjusted time to climb (in minutes) as functions of takeoff weight and target altitude. Note the ragged nature of the surface in the upper panel, which results from table entries being rounded to the nearest minute, and that in the lower panel most of this discretization has been removed.	111
E-3	Fuel to climb (in klb) as a function of takeoff weight and target altitude.	112
E-4	Distance to climb (in nautical miles) as a function of takeoff weight and target altitude.	112
E-5	Time to climb as a function of target altitude for a takeoff weight of 560 klb. The dots represent values from the performance manual, while the gray line represents the seventh-order least-squares fit used to remove the discretization from the tabulated numbers. Note that this fitted line was only used to recreate values of climb time at the tabulated values of target altitude: the “hump” at about 4 kft does not interfere with the model.	113
E-6	Cruise fuel flow (in thousands of pounds per hour) and true airspeed (in knots) as functions of weight and altitude.	114
E-7	Time, fuel, and distance to descend, each as a function of initial altitude. Speeds: M0.88/340kt/250kt.	115
E-8	True airspeed during the descent, as a function of altitude. The descent is nominally flown at Mach 0.88, then 340 kt, and finally 250 kt below 10,000 ft. The fine lines are hypothetical curves for Mach 0.88, 340 kt, and 250 kt. The bold line represents the speeds obtained from Equation E.6.	116
E-9	Percentage error in estimated fuel burn, as a function of cruise integration time-interval. Note that the fuel burn for the 3.2×10^{-6} hour interval was used as the baseline for the calculation of percentage error.	116

List of Tables

3.1	Some data for the subject pilots in the first, free-response survey. “FMC experience” refers to flight experience in FMC-equipped aircraft.	31
3.2	Some data for the subject pilots in the second, forced-response survey.	33
3.3	The averages, \bar{R} , of the ranks assigned by the pilots in the second (forced-response) survey, for the six most popular attributes of the outcome.	33
3.4	The averages, \bar{R} , of the ranks assigned by the dispatchers in the second (forced-response) survey, for the six most popular attributes of the outcome.	34
3.5	Average rank, \bar{R} , assigned by the pilots to each of the five sources of information for the selection of a new cruise altitude.	35
3.6	Desired FMS/CDU features mentioned by the 32 subjects, and the frequencies with which they were mentioned.	36
4.1	Aggregate data for the pilots in the utility-assessment experiment.	37
4.2	Attribute data for the pilot subjects in the utility-assessment experiment.	38
4.3	Estimates of the number of points at which utility must be assessed to specify a DM’s preference structure in an n -attribute decision space (i) when a full utility function must be used, and (ii) under the assumptions required for the use of a multilinear utility function.	39
4.4	Least-squares risk aversion by attribute and subject. Negative values indicate risk-seeking behavior.	49
4.5	Mean risk aversion, and p -values associated with a two-tailed null hypothesis that the mean was zero, by attribute and subject.	50
5.1	Parameters used in the GA to produce the results shown in the figures below.	61
5.2	The increase in cost associated with several different vertical structures of the airspace. Note that percentage changes are based on the 1000-ft separation cost.	63
6.1	Tabular data on destination fuel, flight time, and exposure to turbulence, for five candidate trajectories.	70
A.1	Mean fractional disagreement, $\bar{\Delta}$, (averaged over all subjects) between the decisions made by the subjects and those made according to the three models.	83
A.2	Correlation coefficients for linear relationships between objective distance measures and subjective difference between the alternatives for the three models.	83
B.1	The attributes of the outcome mentioned by the 32 subjects, and the frequencies with which they were mentioned.	85
B.3	The rankings of six attributes of the outcome given by the subject pilots. Concordance, $W = 0.701$	89

B.4	The rankings of six attributes of the outcome given by the subject dispatchers. Concordance, $W = 0.748$.	90
B.5	The rankings of six sources of information given by the subject pilots. Concordance, $W = 0.345$.	90
C.1	Multilinear utility coefficients by subject.	97
E.1	Coefficients of the powers of altitude, A , in the models for descent time, fuel, and distance. Also shown are the coefficients of determination, R^2 , for each model. *Starred coefficients were constrained to be zero.	109

Chapter 1

Introduction

1.1 Motivation

This research was primarily motivated by the need at NASA Ames to model more aspects of human decision making than were embodied in the rule-based models of MIDAS¹ (Corker & Smith, 1993). The research charter was to improve models of value-based decision making, and to apply them to appropriate areas of the pilot's task. Early in the course of the research, it became clear that most aeronautical decisions are not value-based in the decision theorist's sense—although they may have been when aviation was less regulated—but rather are rule-based procedures, or follow some other simple paradigm (Patrick, 1993).

This realization prompted examination of several different cockpit decision-making problems. The remainder of the research was focused on a particular decision problem—one which was both hypothesized and observed to involve value-based decision making, and for which there appears to be a significant need for decision support systems: the optimal constrained vertical navigation of long-haul aircraft (Patrick, 1995).

1.2 Objectives

Broadly, there are two aspects to any optimization problem: deciding what criterion constitutes optimality, and finding solutions which are optimal with respect to this criterion. This thesis considers both in the context of the vertical navigation of long-haul aircraft. The main objectives of this research may be divided into the three areas shown in Figure 1-1.

The Engineering Objective is to improve the fuel consumption, safety, economy, schedule adherence, ride comfort, and other attributes of a flight, by providing the decision makers with a suitable decision support system, or an automatic trajectory optimization system, as appropriate.

The Scientific Objectives are (i) to understand where pilots make value-based decisions, (ii) to improve the understanding of the way in which human pilots make decisions involving trajectory selection; and (iii) to examine the implications of this decision-making behavior on air transportation system design.

¹NASA's Man-Machine Integration Design and Analysis System




Figure 1-1: The engineer's hierarchy of research objectives, in which the engineering problem drives both the scientific questions and the experimental objectives. Note that a physicist might reverse the precedence of the engineering and science items.

The Experimental Objectives are (i) to determine which attributes of the outcome of a flight are important to pilots, and how these attributes differ from those which are important to dispatchers; (ii) to model pilots making decisions involving trade-offs between these attributes, and evaluate how such a model might be different from that used by the airlines; and (iii) to build an experimental system to explore trajectory optimization under a wide set of constraints and objectives.

1.3 The Profile Selection Problem

1.3.1 Decision Makers

Aircraft guidance decisions for an air carrier flight—like the typical route shown in Figure 1-2 from Los Angeles to Sydney, Australia²—are made by three main groups: the pilots (there are typically two), the airline's operational control center (AOC, or *dispatch*), and Air Traffic Control (ATC), as shown in Figure 1-3. Currently, the airline's AOC plans the flight (Grandeau, 1995; Beatty, 1995), and produces a flight plan, which the pilot—who is the final authority in matters affecting the safety of the aircraft—can accept or revise. One of the pilot's jobs is to take wind and weather into account in-flight, and potentially make altitude changes enroute after coordinating with AOC and ATC. Much work has been done on the group decision-making involved in this process (e.g. Orasanu et al., 1993; Smith et al., 1994), but value-based models of the individual decision-maker's behavior are not well developed in this context.

The advent of *Free Flight* (RTCA, 1994)³ may change some of this decision-making process, by removing ATC from their current role as providers of enroute separation between aircraft. If and when this change occurs, there will be increased scope for the kinds of path-selection decision making considered in this research.

1.3.2 Search Space

The decision-making problem is complicated by the number and complexity of the elements facing the pilots. Some of these are described below.

²Coastline data for Figure 1-2 were obtained from the U.S. Geological Survey (1996).

³Formerly the Radio Technical Commission for Aeronautics.




Figure 1-2: Los Angeles to Sydney: a typical trans-Pacific route.




Figure 1-3: The decision-making trinity in air carrier operations.

Vertical Structure of the Airspace

Federal Regulations (see Federal Aviation Administration, 1996b, FAR 91.179.b) require that pilots flying under Instrument Flight Rules (IFR) choose from a relatively limited set of cruise altitudes: below 29 kft MSL⁴ at thousand-foot intervals (e.g. westbound at 12 kft, eastbound at 13 kft, westbound at 14 kft and so on); and above FL 290 at spacings of 2 kft (e.g. east-bound at FL 290, westbound at FL 310, eastbound at FL 330 and so on).⁵ This structure is shown in Figure 1-4. Some of these restrictions may be lifted in the future as the accuracy of modern altimetry is recognized, and more are being lifted under the *National Route Program*⁶ or may be lifted under the *Free Flight* initiative.

Assuming only vertical exploration, no lateral exploration, a flight composed of s segments, each of which can be flown at any of l flight levels will have $n = l^s$ unique profiles. Consider the flight shown in Figure 1-2, which is divided into fifteen segments, each of which might be flown at one of perhaps 10 altitudes: there are 10^{15} possible profiles in the search space. Clearly, the problem of choosing a path in this exponential search space is not trivial.

⁴Above mean sea level.

⁵For example, an eastbound pilot may typically only choose from the following altitudes (in thousands of feet, or kft): 5, 7, 9, 11, 13, 15, 17, 19, 21, 23, 25, 27, 29, 33, 37, 41, 45, 49, 53.

⁶The *National Route Program* is a very limited version of free flight which is currently operating.




Figure 1-4: The present vertical structure of U.S. Airspace. Note that below 29 kft opposing traffic is separated by one thousand feet, whereas above 29 kft it is separated by two thousand feet. Class A (IFR-only) airspace extends from 18,000 feet MSL to 60,000 feet MSL.

Aircraft Performance Variation

Also, as can be seen from Figure 1-5, the performance of a typical transport aircraft varies significantly within its weight-altitude operating envelope. This makes the job of predicting the performance consequences of a flight plan too difficult for the human pilot. Since fuel accounts for between 10 and 15% of an airline's expenses (Rubbert, 1994; Trujillo, 1996), planning a flight to minimize operating cost is important.

Weather

Weather produces many factors which affect flight planning: turbulence, icing, thunderstorms, and winds, to name a few. Some of these factors act as constraints on solution trajectories, but the majority are better thought of as negative factors in the outcome: factors which can be tolerated to some extent, but whose effects should be minimized. Timely and accurate dissemination of weather information, particularly winds and temperatures aloft, although very important for accurate flight planning (Barrows, 1993), will not be addressed in this thesis.

1.4 Outline

In Chapter 2, a broad framework for decision modelling is proposed. Within this framework, many models of decision making—both prescriptive and behavioral—can be considered. Several behavioral models of human decision making are discussed. Notions of optimality are then considered, starting with the way in which optimality is currently defined by the airlines, and implemented in the typical FMS, and concluding with a discussion of the prescriptive and behavioral uses of utility theory, and its limitations. In Chapter 4, an experimental determination of a utility-maximization model




Figure 1-5: Fuel economy (in nautical air miles per klb of fuel) against aircraft weight and altitude showing the large variation in cruise performance over the aircraft's operating envelope. Data are at Long-Range Cruise (LRC) for the Boeing 747-400 with PW4056 engines.

of pilot decision making in the context of trajectory selection is described, and is compared with several other behavioral models.

In Chapter 5, algorithms for finding optimal solutions for the trajectory selection problem are considered. One of these, the Genetic Algorithm or GA, is developed, and its many desirable properties are enumerated: it can, for example, handle difficult-to-model constraints, and a wide variety of objective functions like the non-linear utility functions elicited in Chapter 4. Finally, in Chapter 6, the two main themes of optimality and optimization are combined into a decision aid, which is proposed to help pilots with the trajectory selection problem.

◇ ◇ ◇

Chapter 2

Decision Models

Prescriptive and Behavioral

In this chapter, the nature of decision making is considered, several *prescriptive*¹ and *behavioral*² decision models are explained, and one—utility maximization—is developed. Its axiomatic basis is described in detail, and is then used to draw conclusions about the efficacy of modelling different types of aeronautical decisions.

2.1 A Framework for Decision Modelling

It is instructive to think of decision making in terms of the three-stage framework of Figure 2-1: (1) *selection* of some or all of the alternative courses of action in the *search space*,³ (2) *evaluation* of their consequences, which maps the points in the search space into points in a *decision space*, and (3) *comparison* of the points in the decision space, using a multi-attribute value or utility function (or some simpler criterion) to produce points on a *value or utility scale*. Any or all of these stages may be executed once or iteratively, allowing this framework to encompass a diverse set of decision models.

From a prescriptive point of view the first two tasks are straightforward: find all allowable trajectories,⁴ and evaluate each of them using a high-fidelity flight simulation.⁵ A prescriptive model for the third task, however, is harder to define. From a behavioral point of view, while the search and evaluation aspects of the larger decision problem are rich areas for research, it is again the third task with which this thesis is initially concerned.

Perhaps only the third task in Figure 2-1, *comparison*, would be thought of as decision making by a traditional decision theorist, but it is difficult to compare models in the aviation domain without considering a broader definition of decision making.

¹A *prescriptive* model defines behavior which is optimal with respect to some objective criterion. I have a slight preference for the term *prescriptive* (which means based on prescription) over the term *normative* (which means of or establishing a standard or customary behavior—which may or may not be optimal).

²A *behavioral* model is one which describes human behavior, which may or may not be optimal according to some objective criterion.

³The *search space* consists of either the physical dimensions of space and time in which courses of action are executed, or some other representation in which the actions can be described.

⁴In practice, it can be hard to determine which trajectories are allowable according to some constraints (e.g. fuel burn) without first performing the evaluation step.

⁵This relies on the assumption that search and evaluation costs are negligible when compared to the opportunity cost associated with making a suboptimal decision.




Figure 2-1: A decision making paradigm with three operations: *selecting* admissible solutions (represented here by gray dots) in the *search space* ($s_1 \dots s_3$), *evaluating* each candidate to determine its attribute location (x, y) in the *decision space*, and finally *comparing* the candidate solutions, by mapping them all onto a single *value* or *utility scale* (v).

2.2 Simple Decision Models

2.2.1 Behavioral Models

In the decidedly behavioral field of Naturalistic Decision Making (NDM) humans are modelled as making satisficing (Simon, 1955) rather than optimizing decisions: instead of looking for the best course of action (COA), they may execute the first acceptable COA they find. This corresponds to selecting a single point in the search space of Figure 2-1, evaluating it, and—if it produces acceptable outcomes—selecting it by default. Klein (1993) has observed such recognition-primed decision making in fireground commanders and military personnel—who assess a situation, and then simply imagine or recall an apparently suitable course of action, run through its consequences in their minds, and—if it is appropriate—execute it. There are many cockpit decisions for which this model seems appropriate: choosing a point to turn from the downwind leg to the base leg before landing at an unfamiliar airport, for example. Such a decision model does not qualify as prescriptive, however, unless the time pressures or mental effort involved in a more optimizing approach are considered, and would be prohibitive.⁶ Lacking the third stage—*comparison*—such decisions are not particularly fertile ground for efforts at value-based modelling.

Lexicographic ordering (de Neufville, 1990) makes some use of attribute values, while remaining simple enough to have some behavioral application to the final operation of Figure 2-1. In this scheme, the subject orders all relevant attributes in the decision space according to their importance, then ranks each alternative according to its score on the most important attribute. The second attribute is used to break any ties, and then the third is used to break any remaining ties, and so on.⁷

The Analytic Hierarchy Process (Saaty, 1977, 1990), in which intransitivities in preference are resolved in the principal eigenvector of a *judgment* matrix, does not seem to embody either a prescriptive or a behavioral model of decision making. It is at best a method for analyzing subjective data (Yang & Hansman, 1995), and is considered by some (Dyer, 1990) to produce rankings which

⁶In such a case, a decision aid might profitably be deployed to help the human operator make a more optimal decision.

⁷Asimov (1970) offers an interesting—though unworkable—example of lexicographic ordering in his *Three Laws of Robotics*.

are too arbitrary even for this purpose.⁸

2.2.2 Proceduralized Decision Making

Many aeronautical decisions have been proceduralized, or turned into standard procedures. For example, pilots are extensively trained to respond to a wide variety of emergency situations with an appropriate predetermined sequence of actions—a procedure. This kind of decision-making behavior bypasses all but the first stage of the model of Figure 2-1, and is therefore not amenable to value-based modelling.

2.2.3 Prescriptive Models

The Airline's Notion of Optimality—What the FMS minimizes

Some optimization is currently performed by the FMSs installed in aircraft such as the Boeing 747-400 (Honeywell, 1994) and the Boeing 737-300 (Schreur, 1995). The objective function embodied in these FMSs is the total monetary cost of the flight, C , which is defined in terms of fuel burn, f , in thousands of pounds, and flight time, t , in hours:

$$C = 10 c_f f + c_t t \quad \text{dollars,} \quad (2.1)$$

in which c_t is the operating cost (excluding fuel) in dollars per hour, and c_f is the cost of fuel in cents per pound.⁹ By defining a cost index, I_c , as the ratio of these two costs, c_t and c_f , Equation 2.1 may be rewritten as:

$$C' = f + \frac{c_t}{10 \cdot c_f} t = f + \frac{I_c}{10} t. \quad (2.2)$$

Because C and C' differ by only a constant factor, any policy (i.e. flight plan) which minimizes one, minimizes the other. Operationally, the pilot enters a cost index determined beforehand by the airline, and when it is in the ECON mode,¹⁰ the FMS acts to minimize C' .

Clearly, airlines use the FMS to minimize a *linear* combination of fuel and time. This institutional cost function may be appropriate at the higher levels of airline management, where fuel consumption is the sum of fuel consumptions on many thousands of flight each year. There are, however, two ways in which this objective function might be considered inadequate. First, it does not take into account the fact that, after takeoff, resources like fuel and time are constrained. Second, it spans only two attributes.

Constrained Resources

The form of the pilot's cost function for fuel and time may be quite different: studies of human decision-making in other fields (e.g., blood-bank management in Keeney & Raiffa, 1993) suggest that an individual's preferences are noticeably non-linear.

Before an aircraft leaves the ground, fuel can be added or off-loaded and the departure time can be adjusted to meet all of the foreseeable constraints imposed on the flight. These resources are *tradeable*: there is an unlimited supply of fuel, available at a constant price. To maximize its profit from providing a large number of revenue-generating flights, an airline needs to be most concerned with cost. Safety, and therefore fuel reserves are not part of the airline's objective function: they

⁸It seems roughly to produce the cardinal equivalent of the ordinal Kendall coefficient of concordance used in Section B.2.

⁹The factor of 10 is required to convert the fuel cost, c_f , from cents per pound of fuel to dollars per thousand pounds of fuel.

¹⁰Economy mode.

are treated as constraints. However, outcomes of a flight—such as fuel and time—are not tradeable once an aircraft is airborne. Thus there is no market within which their prices can be set.¹¹

Given the multidimensional nature of both the decision problem and the pilot’s values, and given the constraints on several of the resources involved, it is appropriate to turn to utility theory for a quasi-prescriptive model of decision making behavior.

2.3 Utility Theory

2.3.1 Axiomatic Basis

Utility theory is built on six axioms (de Neufville, 1990; Keeney & Raiffa, 1993). The first three form a basis for the construction of a *value* function, $\mathcal{V}(x)$, and the latter three provide the necessary extensions for the construction of a *utility* function, $\mathcal{U}(x)$.

Complete Preorder of Preferences For all possible pairs of outcomes, x_a and x_b , the subject either prefers one, or is indifferent to them: $x_a \succ x_b$, $x_a \prec x_b$, or $x_a \sim x_b$.

Transitivity of Preferences For all possible sets of outcomes, x_a , x_b , and x_c , if $x_a \succ x_b$ and $x_b \succ x_c$, then $x_a \succ x_c$.

Monotonicity of Preferences For every outcome, x , between x_0 and x_1 , where $x_0 \prec x \prec x_1$, there exists a number w , where $0 \leq w \leq 1$, such that $\mathcal{V}(x) = w \mathcal{V}(x_1) + (1 - w) \mathcal{V}(x_0)$.

Existence of Probabilities Probabilities exist, and can be measured.

Monotonicity of Probabilities For a lottery with a fixed pair of outcomes, increasing the probability of the more desirable outcome increases the desirability of the lottery: for all x_a and x_b , where $x_a \succ x_b$, if $p_1 > p_2$ then $\langle x_a, p_1; x_b \rangle \succ \langle x_a, p_2; x_b \rangle$.

Substitution Preferences are linear in probability: for all x_a , x_b , and x_c , if $x_a \sim x_b$ then $\langle x_a, p; x_c \rangle \sim \langle x_b, p; x_c \rangle$.

That human subjects might behave according to the first three axioms and the fifth axiom is not hard to postulate. There is, however, some debate about the applicability of the sixth axiom to human decision making.

2.3.2 Model Decisions About Performance, Not Safety

It is much easier to define a better outcome in domains which are not safety-related. While it is not unethical to trade off economic cost against risk of injury or death (Keeney, 1995), it is very difficult. Two postulated utility functions are shown in Figure 2-2. On the left is a utility function for the outcome of a safety-related task such as landing an aircraft, where x might represent the distance between the centerline of the runway and the aircraft’s touch-down point. Note that the worth of the outcome is relatively insensitive to changes in the outcome, except at the extremes, where in this example the aircraft touches down beside the runway, usually with disastrous results. On the right is the utility function for a performance-related task, in which x might represent effort expended on a project. In this case, the value of the effort increases fairly steadily, and there are no “pathological” outcomes with very (unmeasurably) low values. There are two problems with assessing value for the safety-related task on the left. First, the bad regions are extremely bad: if there is a utility

¹¹It is true that fuel and flight time can—to some extent—be traded against one another *within* a single flight. However, this does not constitute a market, which involves the trade of a single type of good.

function which maps x into the standard utility interval of $[0, 1]$, then these extrema dominate, and everywhere else the utility is indistinguishable from 1. Second, the outcomes associated with low value for the safety-related task are—either by design or naturally—very unlikely events. Examples of such unlikely events include in-flight structural failure or engine failure.




Figure 2-2: Hypothetical examples of two types of utility, $U(x)$, as functions of outcome, x , (a) for a safety-related task, and (b) for a performance-related task. Note the avoidance states at one extreme of x in (a).

Extreme Values

There are outcomes in aviation whose values are so extreme as to be unmeasurable. While actuaries routinely calculate the value of a life, an individual decision maker cannot be expected to appreciate the value of his own—even though it would seem to be a prerequisite for rational decision making in the face of threats to his safety, and aviation is full of these.

Extreme Probabilities

The fourth axiom also presents a problem. Although it seems reasonable that probabilities exist, it is not clear that they can be measured easily—much less comprehended by human decision makers—for outcomes which happen only very rarely. Even where accident rates can be measured, there is often insufficient experience to quantify the associated probabilities precisely (Patrick, 1996).

Previous experiments conducted by the author (See Appendix A) involved observing pilot's decision-making in situations in which the safety of the outcome was sometimes in doubt. Not only do most safety-related decisions in aviation follow the simple decision paradigms of Section 2.2, but those which might be value based are not amenable to modelling with utility theory for the reasons given above.

2.3.3 Prescriptive Use of Utility Theory

Decisions like those involved in altitude selection are often made by people, because avionics designers do not have objective functions with which to automate the decision-making. But without such an objective function, there is no prescriptive model against which to evaluate the optimality of the human's decisions. In other words, the situations in which a human is the decision maker are usually those in which it is impossible to judge the quality of his decisions. This has been called Roseborough's dilemma (Roseborough, 1988).

Human decisions are often conservatively judged against a cost-minimizing (e.g. expected monetary value, or EMV) model. An alternative is to use the notion that—to the extent that he is

consistent—the human decision maker (DM) always makes the correct decision. A sensible middle ground is to measure the DM’s consistency with decisions made according to a model of his preferences: the subject’s utility function is derived after observation of the subject’s choices, and the subject’s decisions are then judged against this model. This might be called the *utility-consistency heuristic* of behavioral decision theory, and it is this criterion against which the subjects observed in Chapter 4 were measured.

◇ ◇ ◇

Chapter 3

Pilot and Dispatcher Surveys

This chapter describes two surveys given in the course of this research. The first, an exploratory survey given only to pilots, was used to determine which attributes constituted the decision space for the vertical navigation problem, and to discover where pilots obtained the information they used in making their decisions. The second, given to both pilots and dispatchers, was used to determine both groups' rankings of the decision-space attributes.

3.1 An Exploratory Survey of Pilots

no. of subjects	32
mean age	48.7 yrs
mean time flying	27.8 yrs
mean flight experience	15,600 hrs
mean FMC experience	4,600 hrs
fraction who were captains	56%

Table 3.1: Some data for the subject pilots in the first, free-response survey.
“FMC experience” refers to flight experience in FMC-equipped aircraft.

In an effort to determine which attributes of the outcome of a flight are considered important, thirty-two line pilots—all of whom were captains or first officers flying for major U.S. air carriers (see the data in Table 3.1)—were asked to list any elements of the outcome of a flight which they felt were important:

Question 1 “In evaluating a flight plan, what elements of the outcome of a flight are important to you? Please list as many relevant elements as you can think of.”

Their answers are shown below in Figure 3-1, which is constructed from the raw data in Table B.1. Time and weather responses have been aggregated into super-categories for comparison, since it was hard to distinguish exactly what the subjects had meant by each response. If specific categories had been provided for responses to the question, then the classification of the responses would have been much easier, but many unexpected responses might never have been given by the subjects.

Note the distinction between *flight time*, which is the duration of the flight, and *schedule adherence*, which is the extent to which the flight arrives on time. A similar distinction exists between fuel burn and fuel at destination. For both of these pairs, such a distinction is possible since there are two controllable variables: in the case of fuel, both fuel burn and destination can be controlled separately by adjusting the takeoff fuel load and the aircraft's route.

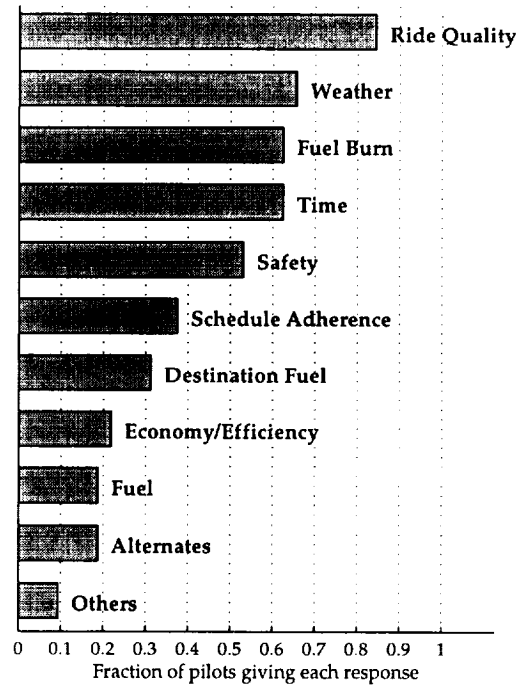


Figure 3-1: The attributes of the outcome mentioned in the first (free-response) questionnaire by the 32 pilots, and the fraction of pilots who mentioned each one.

Surprisingly, pilots mentioned ride quality more often than any other single attribute—even more often than they mentioned safety. This observation helped focus this research on modelling the decision making involved in avoiding turbulence, and trading off turbulence exposure against fuel burn and flight time.

3.2 A Forced-Response Survey of Pilots and Dispatchers

A problem with the free-response format of this questionnaire was that pilots were free to mention items that aren't attributes in the decision space, but rather are part of the search space.¹ Weather is a good example of this: inclement weather may force a detour, resulting in a longer flight with a greater fuel burn; it may contribute to increased turbulence, thus lowering ride quality; or it may increase the probability of an accident. However, weather shouldn't be accounted for separately from its effects—unless it produces a constraint on operations. Its value should be measured by its effect on attributes such as flight time and fuel burn.

3.2.1 Pilots' Rankings of Attributes of the Outcome

To make up for this shortcoming of the free-response format, but armed with the wide range of responses it produced, a second questionnaire was devised. In this second survey, thirty five pilots (28 of 32 respondents from the first questionnaire, and 7 additional subjects—see the data in Table 3.2) were asked the following question.

¹Recall the specific definitions of the search space and the decision space in Section 2.1.

no. of subjects	35
mean age	49.3 yrs
mean time flying	28.8 yrs
mean light experience	15,600 hrs
mean FMC experience	4,400 hrs

Table 3.2: Some data for the subject pilots in the second, forced-response survey.

Question 2 “How important are these elements of the outcome of a flight to you? Please mark the importance of the following six items (flight time, schedule adherence, fuel burn, fuel at destination,² ride quality—i.e. passenger comfort, and safety) on the scale below. Your ordering of the items is most important.”

The position of each item along the scale, which had been marked from “irrelevant” to “very important”, was converted into a rank.³ The raw results are shown in Table B.3, and the rankings are summarized here in Table 3.3.

<i>Attribute</i>	\bar{R}
Safety	1
Ride quality	2.57
Fuel at destination	3.44
Schedule adherence	4.14
Fuel burn	4.33
Flight time	5.51

Table 3.3: The averages, \bar{R} , of the ranks assigned by the **pilots** in the second (forced-response) survey, for the six most popular attributes of the outcome.

In fact, it is not possible to separate some of these attributes. Fuel burn and fuel at destination are linked: they must add to the fuel load at takeoff. Flight time and time ahead of schedule must sum to the difference between actual departure time and scheduled arrival time. For this reason these attributes were combined into only two: fuel at destination, and flight time. Given that safety is an attribute which we wish to avoid modelling—following the logic expressed in Chapter 2—this leaves only three attributes in the decision space: fuel at destination, flight time, and ride quality.

A Measure of Agreement Between Judges The Kendall coefficient of concordance, W (see Seigel & Castellan, 1988, p. 262) was used as a measure of the agreement between the $k = 35$ pilots’ (judges’) rankings of the $n = 6$ attributes shown in Table B.3:

$$W = \frac{\sum_{i=1}^n (\bar{R}_i - \bar{\bar{R}})^2}{n(n^2 - 1)/12}, \quad (3.1)$$

²“Fuel at destination” refers to the quantity of fuel remaining in an aircraft’s tanks upon landing at the destination airport, rather than to the availability of fuel at that airport.

³Although the locations of the items could have been measured to provide a more cardinal estimate of importance, it was felt that such a measure would be unreliable. Respondents varied in their treatments of the scale—for instance by marking all the items near one or the other of the extremes—sufficiently to bring into question the cardinality of the data. For this reason, only the ordinal information (i.e. rank) was used.

where \bar{R}_i are the mean rankings for each of the n attributes:

$$\bar{R}_i = \frac{1}{k} \cdot \sum_{j=1}^k R_{ij}, \quad (3.2)$$

and $\bar{\bar{R}}$ is the grand mean, or mean of the \bar{R}_i 's. W can lie anywhere on the range $[0, 1]$. If W has a value of 0, then there is complete disagreement between the judges' rankings, whereas if W has a value of 1, then there is complete agreement between the judges' rankings. Values between 0 and 1 imply partial agreement.

For the data from the pilots in Table B.3, $W = 0.701$. Since the critical value⁴ of W for a 1% level of confidence, $W_{1\%}$, is less than 0.146, this indicates strong, statistically significant agreement among the pilots.

It is no surprise that safety is ranked first by all of the subjects. In fact, any other rank for this item would be a cause for concern! It is therefore illuminating to remove safety from the table, and recompute the level of concordance. In this case, with $n = 5$ attributes, $W = 0.477$, which is still well above the critical value: $W_{1\%} \leq 0.16$.

3.2.2 Dispatchers' Rankings of Attributes of the Outcome

The same question (Question 2) was put to dispatchers at a major U.S. airline's operational control center (or AOC). The means of the ranks assigned by the dispatchers are shown in Table 3.4, while the raw data are shown in Table B.4.

<i>Attribute</i>	\bar{R}
Safety	1
Ride quality	2.57
Schedule adherence	2.86
Fuel burn	4.71
Fuel at destination	4.86
Flight time	5

Table 3.4: The averages, \bar{R} , of the ranks assigned by the **dispatchers** in the second (forced-response) survey, for the six most popular attributes of the outcome.

For the dispatchers' ranks, $W = 0.748$. Since the critical value of W for a 1% level of confidence is 0.398, this indicates strong agreement between the dispatchers. After removing safety from the rankings, as for the pilots results, W dropped to 0.559, which is still statistically significant.

3.2.3 Agreement Between Pilots and Dispatchers

Figure 3-2 shows the mean rankings produced by both groups—pilots and dispatchers—graphically. If all the points had been on the gray line, there would have been perfect agreement between the two groups' rankings. There is, however, some disagreement. It is interesting to note that this disagreement involves destination fuel, which pilots seem to value more highly than dispatchers do, and schedule adherence, which dispatchers seem to value more highly than pilots do. There is, however, complete agreement between the two groups that safety and ride quality have the highest priorities among these six items.

⁴Critical values of W for $n = 6$ are only tabulated for $k \leq 20$ in Seigel and Castellan (1988). Since these values diminish as k increases, the value for $k = 20$ is taken as a conservative estimate of the value for $k = 35$.




Figure 3-2: A quasi-cardinal representation of the mean ranks assigned by both pilots and dispatchers to each of the six attributes of the outcome.

3.2.4 Pilots' Rankings of Sources of Information

In order to determine whence pilots obtain the information they require for decision making about vertical navigation, thirty pilots (again 28 of the 32 respondents from the first questionnaire, and this time 2 additional subjects) were asked to rank 5 sources of information⁵ for the selection of a new cruise altitude in the following question.

Question 3 “How important are the following sources of information to you when selecting a new cruise altitude? Please mark the importance of the following five items—Flight Management System (FMS), Flight Plan, Dispatch, Pilot Reports (PIREPS), Pilot’s Operating Handbook or Flight Manual (POH)—on the scale below. Your ordering of the items is most important.”

The location of each item along a linear scale (again marked from “irrelevant” to “very important”) was used to determine the rank of the importances of each item, from 1 (most important) through 5 (least important). These results are shown in full in Table B.5. Average rank for each of the i items was calculated using Equation 3.2, with $k = 30$. The raw data are presented in Table B.5, and are summarized here in Table 3.5.

<i>Source of Information</i>	\bar{R}
PIREPS	1.52
FMS	2.88
flight plan	3.13
POH	3.47
dispatch	4.00

Table 3.5: Average rank, \bar{R} , assigned by the pilots to each of the five sources of information for the selection of a new cruise altitude.

⁵In the initial, exploratory survey, the pilots had been asked to list sources of information which they considered important. The five most frequently mentioned responses were used in this question in the second survey.

For the thirty pilots' rankings of the five sources of information, $W = 0.345$. Since the critical value⁶ of W for a 1% level of confidence is less than 0.160, this again indicates good agreement between the pilots.

Most noticeable in these data is the fact that pilots place *dispatch*—a body of people who are trained and certified to help pilots in such decision-making tasks—last in the list of sources of information for altitude planning.

3.3 FMS Features Desired by Pilots

As part of the first survey, the pilots were also asked to describe any additional features they would like in their FMSS to help them with altitude selection and flight planning.

Question 4 “If you could have any additional features in the FMS to help with these tasks, what would you want?”

The individual responses to this question are listed in the long table in Section B.1.1, along with the principal aircraft the pilots were flying at the time of the survey. The responses and frequencies with which they were given are summarized in Table 3.6. The two most frequently requested

<i>Function</i>	<i>frequency</i>
Improved display of...	16
performance	7
weather and winds	3
terrain	3
airspace	2
traffic	1
Improved flight planning...	11
vertical	7
unspecified	3
horizontal	1
Better use of wind data...	9
uplink of data	7
calculation of effects	2
display	1
Nothing	6
Improved performance model	1

Table 3.6: Desired FMS/CDU features mentioned by the 32 subjects, and the frequencies with which they were mentioned.

features were improved display of performance, and improved vertical flight-planning capability. This is perhaps not surprising, since the question was asked in the context of vertical navigation decision making. What is noteworthy, however, is that only 6 of the 32 subjects desired no additional features.

◇ ◇ ◇

⁶Again, the critical value of W ($n = 5$, $k = 20$) of 0.160 is used as a conservative estimate of the true 1% value.

Chapter 4

Utility Assessment Experiment

Eliciting Pilots' Objective Functions

In this chapter, one particular model of decision making for the trajectory selection problem—maximization of expected utility—is constructed by experiment, for the attributes determined in Chapter 3, and validated by comparison with several of the other decision models described in Chapter 2.

4.1 Background

Active line pilots—all of whom were flying or had recently flown intercontinental routes for major U.S. air carriers—were interviewed at NASA Ames Research Center. Some aggregate data for the group are shown in Table 4.1. The pilots were each asked questions (1) to determine which attributes they considered relevant, (2) to elicit their utility functions over those attributes, and (3) to make many comparisons between pairs of points (i.e. attribute combinations) in the decision-space.

no. of subjects	5
mean age	51.2 yrs
mean time flying	30.2 yrs
mean flight experience	14,800 hrs
mean FMC experience	2,520 hrs

Table 4.1: Aggregate data for the pilots in the utility-assessment experiment.

The subjects were instructed to answer the questions as though they were planning a long intercontinental flight on a route with which they were very familiar (e.g. Los Angeles to Sydney, Australia), in the aircraft they would normally fly on that route. The weather, except for winds and turbulence, was described as no factor, and the cost index (as used in Equation 2.2) was specified as $I_c = 100$. Subjects were also informed that there were no “correct” answers to any of the questions they would be asked: their answers would only be judged by their consistency.

4.2 Verification and Scaling of Important Attributes

The subjects were asked to confirm that the first three independent attributes determined in Section 3.2—destination fuel, flight time, and ride quality—were the most important decision attributes

of a flight, which they all did. The pilots were then asked to set end points for destination fuel and flight time at the smallest and largest values each pilot considered reasonable for the route he had chosen.

Ride quality was a more difficult attribute to deal with, since there exists no universally accepted measure which covers both its strength and duration.¹ The subjects were asked questions to determine what level of turbulence² they would endure for between zero and 60 minutes to save 2000 lb of non-critical³ fuel. All but one of the subjects, it was determined, would make such decisions for moderate turbulence, the exception would make them for light turbulence. The results of these assessments are presented in Table 4.2, in which f_0 and f_1 are respectively the least desirable (smallest) and most desirable (largest) amounts of destination fuel, and t_0 and t_1 are the least desirable (longest) and the most desirable (shortest) flight times, respectively.

	<i>Aircraft type</i>	<i>from</i>	<i>to</i>	f_0 (klb)	f_1 (klb)	t_0 (hrs)	t_1 (hrs)	<i>turb.</i>
S_1	B-747-400	KLAX	Sydney	30	50	15.0	13.0	mod.
S_2	MD-11	PANC	Seoul	18	30	9.5	7.5	mod.
S_3	B-767	KSFO	London	10	40	10.5	9.0	mod.
S_4	B-747-400	KSFO	Hong Kong	17	37	14.5	12.5	light
S_5	B-747-400	KLAX	Sydney	25	40	16.5	13.5	mod.

Table 4.2: Attribute data for the pilot subjects in the utility-assessment experiment.

4.3 Comparisons of Pairs of Alternatives

The subjects were asked to choose the more attractive alternative from each of forty pairs of trajectories⁴ on the basis of destination fuel, flight time, and length of exposure to turbulence. The data for a typical pair of trajectories are shown below.⁵

[24 klb, 14:00 hr, 20 min] or [30 klb, 14:30 hr, 10 min].

4.4 Using a Multilinear Utility Function

It would be possible to build a full n -dimensional utility function, one point at a time. Assuming that utility over each of the n axes is defined by 5 points, and that the two extrema $[x_0, y_0, z_0]$ and $[x_1, y_1, z_1]$ have utilities of 0 and 1 by definition, there are $5^n - 2$ points to determine, which in this 3-dimensional case would be 123 points. Since each assessment typically takes about half a dozen questions, this process would be very time consuming.

Fortunately, with one important assumption, it is possible to construct a simple alternative model—the *multilinear* utility function—which is equivalent. The required assumption is that the

¹Although there are obvious measures for each element individually, e.g. *light* turbulence, and 20 minutes, respectively.

²The Aeronautical Information Manual (or AIM Federal Aviation Administration, 1996a) defines four levels of turbulence: light, moderate, severe, and extreme.

³Non-critical fuel was defined as fuel which would not be required to reach the destination, or meet minimum fuel requirements upon arrival.

⁴Although these questions were asked in the middle of the experiment, they are described here because their results were used extensively in the following sections.

⁵Subject 4.

n attributes being incorporated into the decision model are utility independent.⁶ In other words if, for all possible combinations of m of the n attributes (where $m \subset n$), the subject's preference among lotteries involving those m attributes is unaffected by the levels of the remaining $n - m$ attributes, then a *strategically equivalent*⁷ multilinear utility function may be constructed for the n attributes.

Given that there was demonstrable utility independence between all subsets of the attributes (as described later in Section 4.7), it was possible to use such a multilinear utility function to model each subject's preference structure in the 3-dimensional decision space of destination fuel, flight time, and turbulence duration. The 3-dimensional incarnation of a multilinear utility function, first proposed in 2-dimensional form by Yntema and Klem (1965) and later generalized by Keeney (1968), is:

$$\begin{aligned} \mathcal{U}_{xyz}(x, y, z) = & k_x \mathcal{U}_x(x) + k_y \mathcal{U}_y(y) + k_z \mathcal{U}_z(z) \\ & + k_{xy} k_x k_y \mathcal{U}_x(x) \mathcal{U}_y(y) \\ & + k_{xz} k_x k_z \mathcal{U}_x(x) \mathcal{U}_z(z) \\ & + k_{yz} k_y k_z \mathcal{U}_y(y) \mathcal{U}_z(z) \\ & + k_{xyz} k_x k_y k_z \mathcal{U}_x(x) \mathcal{U}_y(y) \mathcal{U}_z(z) \end{aligned} \quad (4.1)$$

Note that Equation 4.1 is constructed of 3 conditional utilities, $\mathcal{U}_x(x)$, $\mathcal{U}_y(y)$, and $\mathcal{U}_z(z)$, and 7 scaling constants, $k_x \dots k_{xyz}$. Each of the conditional utilities can again be defined by—say—5 points, but for each attribute the end-point utilities are 0 and 1 by definition, leaving only $3n$ points to cover all n attributes. Of the seven scaling constants, 6 can vary, the seventh is constrained by the fact that $\mathcal{U}_{xyz}(x_1, y_1, z_1) = 1$.

In the general n -dimensional case, there are $2^n - 2$ independent scaling constants to be determined, each of which requires the assessment of a single utility value. Thus the multilinear utility function, where applicable, can be constructed by determining only $2^n - 2 + 3n$ utility points. In this 3-dimensional case, that amounts to 15 points, a substantial reduction from the full 123 points that are required when the independence assumptions do not hold. Table 4.3 outlines the differences between the number of points which must be assessed with and without the ability to use a multilinear utility function.

n	full	multi.
1	3	3
2	23	8
3	123	15
4	623	26
5	3123	45
6	15623	80

Table 4.3: Estimates of the number of points at which utility must be assessed to specify a DM's preference structure in an n -attribute decision space (i) when a full utility function must be used, and (ii) under the assumptions required for the use of a multilinear utility function.

⁶Or if they meet other slightly less restrictive criteria involving partial utility independence and partial preferential independence (see Keeney & Raiffa, 1993, p. 292 for a detailed description of these criteria).

⁷Two utility functions are said to be strategically equivalent if and only if they produce identical decisions between all possible pairs of lotteries.

4.5 Evaluation of the Conditional Utilities

Each subject was asked a series of questions to determine his conditional utility for each of the attributes, one at a time. This involved giving the subject a choice between a lottery,⁸ $L = \langle x_a, 50\%; x_b \rangle$, where $\mathcal{U}(x_a) > \mathcal{U}(x_b)$, and the certain outcome whose utility was being assessed, x_c , as shown in Figure 4-1. If the subject preferred the lottery, then the value of x_c was moved




Figure 4-1: The choice between the lottery $\langle x_a, 50\%; x_b \rangle$ and the certain outcome x_c offered to the subjects.

towards the more desirable outcome x_a , but if the subject preferred the certain outcome, the value of x_c was moved towards the less desirable outcome x_b .⁹ This process was repeated until the subject was indifferent between the lottery and the certain outcome. For this to be true, the utility of the certainty must equal the expected utility of the lottery:¹⁰

$$x_c \sim L \Rightarrow \mathcal{U}(x_c) = \mathcal{U}(L). \quad (4.2)$$

Therefore

$$\mathcal{U}(x_c) = 0.5\mathcal{U}(x_a) + 0.5\mathcal{U}(x_b). \quad (4.3)$$

Since $\mathcal{U}(x_0)$ and $\mathcal{U}(x_1)$ are defined to be 0 and 1 respectively, by starting with the lottery $\langle x_1, 50\%; x_0 \rangle$, $x_{0.5}$ was determined using Equation 4.3. Then, $x_{0.25}$ and $x_{0.75}$ were determined using the lotteries $\langle x_{0.5}, 50\%; x_0 \rangle$ and $\langle x_1, 50\%; x_{0.5} \rangle$ respectively. Finally, the value of $x_{0.5}$ was checked using the lottery $\langle x_{0.75}, 50\%; x_{0.25} \rangle$. If the new value of $x_{0.5}$ did not agree with the first value, then the decision maker was considered not to be acting in a consistent manner, and the conditional utility for that attribute was elicited again.¹¹ Figure 4-2 shows the three component utility functions assessed from a typical subject, along with the exponential functional form chosen to describe them.

⁸The notation $L = \langle x_a, p; x_b \rangle$ denotes a binary lottery, L , in which there is a probability p of obtaining outcome x_a , and a probability $(1 - p)$ of obtaining outcome x_b .

⁹For example, to assess the 0.5 utility point for flight time, between attribute extrema of 12.5 and 14.5 hours, the lottery used would be $L = \langle 12.5, 50\%; 14.5 \rangle$. The subject might first be asked for his preference between L and the certain outcome of 13.5 hours. If he chose L , then the certain outcome was moved towards the more desirable extremum of 12.5 hours, say to 13 hours. On the other hand, if he preferred the certain outcome of 13.5 hours, the value of the certain outcome was moved towards the less desirable extremum of 14.5 hours, say to 14 hours. This process was repeated until the subject was indifferent to—say— L and a certain outcome of 14.25 hours.

¹⁰This relies on the assumption that the subject's preferences are linear in probability: the *Substitution Axiom* discussed in Chapter 2.

¹¹In practice, it is difficult to reassess these indifference values. In this experiment, subjects would remember the values they had given to previous questions and—perhaps to appear consistent—usually gave identical answers when the same question was asked again. Also, since repeated questions do not produce answers which are independent and identically distributed about the a true value, it is difficult to use the standard techniques for statistically analyzing repeated measures to estimate parameters and their confidence limits. These two facts conspire to make the statistical treatment of assessed utilities difficult.




Figure 4-2: One subject's conditional utilities for each of the three attributes, showing the exponential fit of Equation 4.7. (S_4)

4.5.1 The Importance of Resource Constraints

In Chapter 2, the constraints on in-flight resources such as fuel and time were mentioned as important factors in the choice of a normative decision model. To demonstrate human pilots' risk-averse behavior in the presence of such constraints, an extra utility assessment was performed on one subject: his conditional utility for flight time was assessed both with and without a schedule constraint (an operations curfew at the destination airport). Figure 4-3 shows these two utility functions over flight time for the subject. The lower utility function, which is marked with circles, shows the subject's preference over time when there are no constraints on the aircraft's schedule; the upper, marked with solid circles, in the presence of the constraint. Note the concavity of the unconstrained utility: the subject exhibited risk-seeking behavior without the constraint. This was confirmed during the questioning, when the subject mentioned that he was willing to take a gamble for the chance of obtaining the desirable low flight time outcome. The constrained utility function, however, is convex. With a curfew scheduled for 15 hours after departure, the subject was no longer willing to

gamble with the possibility of being late, and having to divert to a nearby airport: he was willing to sacrifice time in the certain outcome to avoid this possibility. Clearly, and not surprisingly, the presence of a constraint has a significant effect on the pilot's decision making. While such behavior is captured in this multilinear utility model, it is not replicated in the FMS's linear cost criterion.




Figure 4-3: One subject's conditional utilities for time, both with and without a schedule constraint at 15 hours. (S_1)

4.6 Fitting the Conditional Utility Functions

Given that $U(x_0) = 0$ and $U(x_1) = 1$ by definition, it is necessary to find a functional form for the conditional utilities which passes through these points, and which provides the best possible fit to the elicited utility points $x_{0.25}$, $x_{0.5}$, and $x_{0.75}$. There are several candidate functions which meet these criteria to varying degrees. Some are shown in Figure 4-4 and are described below.

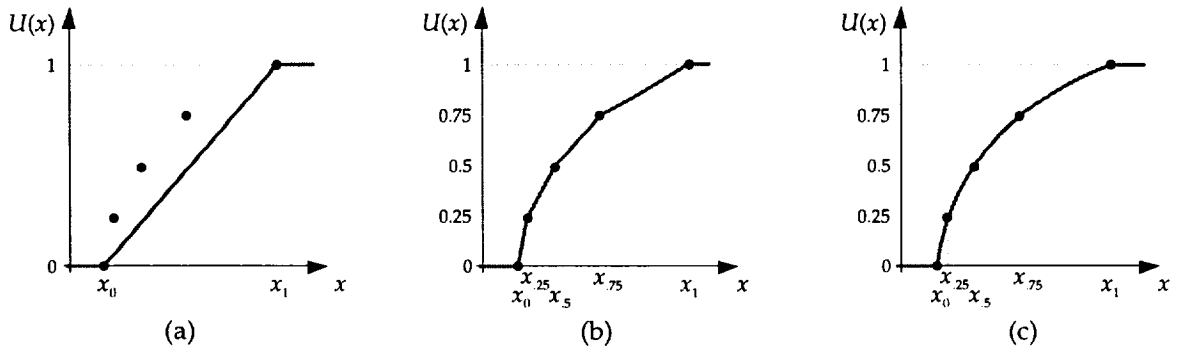


Figure 4-4: Three models for a conditional utility function: (a) linear, (b) piecewise linear interpolation, and (c) exponential.

4.6.1 Linear Utility Function

The use of a linear component utility function corresponds to normalizing the value of an attribute with that attribute's extrema, producing the straight line shown in panel (a) of Figure 4-4:

$$\mathcal{U}_{linear}(x) \equiv \begin{cases} 0 & \text{if } x \preceq x_0, \\ \frac{x-x_0}{x_1-x_0} & \text{if } x_0 \prec x \preceq x_1, \\ 1 & \text{if } x_1 \prec x. \end{cases} \quad (4.4)$$

Lockheed's *Diverter* (Rudolph et al., 1990), a prototypic system for aircraft route planning, uses weights based on rankings of each decision alternative's attainment of each attribute. This is equivalent to using such a linear conditional utility function. Although not appealing as a normative model,¹² the linear conditional utility model of Equation 4.4 is provided for comparison.

4.6.2 Piecewise Linear, or Interpolated Utility Function

The second simple model considered, linear interpolation has the advantage of fitting all of the elicited points by definition, and is shown in panel (b) of Figure 4-4:

$$\mathcal{U}_{interp.}(x) \equiv \begin{cases} 0 & \text{if } x \preceq x_0, \\ 0 + 0.25 \frac{x-x_0}{x_{0.25}-x_0} & \text{if } x_0 \prec x \preceq x_{0.25}, \\ 0.25 + 0.25 \frac{x-x_{0.25}}{x_{0.5}-x_{0.25}} & \text{if } x_{0.25} \prec x \preceq x_{0.5}, \\ 0.5 + 0.25 \frac{x-x_{0.5}}{x_{0.75}-x_{0.5}} & \text{if } x_{0.5} \prec x \preceq x_{0.75}, \\ 0.75 + 0.25 \frac{x-x_{0.75}}{x_1-x_{0.75}} & \text{if } x_{0.75} \prec x \preceq x_1, \\ 1 & \text{if } x_1 \prec x. \end{cases} \quad (4.5)$$

While the most accurate at fitting each of the assessed points, this model is not appealing because it lacks simplicity: in addition to the locations of the attribute extrema, it requires three parameters to specify the DM's preferences, and is susceptible to assessment errors in any one of them.

4.6.3 Least-Squares Exponential Utility Function

A generalized form (capable of handling arbitrary values of x_0 and x_1) of an often-used model for conditional utility (see, for example, Keeney & Raiffa, 1993) is the following exponential function, also shown in panel (c) of Figure 4-4:

$$\mathcal{U}_{exp.}(x) \equiv a - be^{-cx}. \quad (4.6)$$

Writing Equation 4.6 for each of the end points, $\mathcal{U}(x_0) = 0$ and $\mathcal{U}(x_1) = 1$, and solving for a and b gives:

$$\mathcal{U}_{exp.}(x) \equiv \begin{cases} 0 & \text{if } x \preceq x_0, \\ \frac{1-e^{-c(x-x_0)}}{1-e^{-c(x_1-x_0)}} & \text{if } x_0 \prec x \preceq x_1, \\ 1 & \text{if } x_1 \prec x. \end{cases} \quad (4.7)$$

Since the utility function of Equation 4.7 need not pass through all of the elicited points, the single parameter c must be adjusted to minimize the errors in some meaningful way.

¹²When such linear models are combined into a multilinear utility function, they produce a simple surface called a *regulus* in the n -attribute space.

Adjusting the Parameter c Since subjects were adjusting the abscissa (resource variable) values in the experiment, it would not be appropriate to use traditional least squares, as that would involve adjusting errors in the ordinate (utility) estimates. It is more appropriate to perform a least squares minimization of errors in the abscissa values. Following this logic, the parameter c was adjusted to minimize the following measure of error, \mathcal{E} :

$$\mathcal{E} = (x_{0.25} - \hat{x}_{0.25})^2 + (x_{0.5} - \hat{x}_{0.5})^2 + (x_{0.75} - \hat{x}_{0.75})^2, \quad (4.8)$$

where

$$\hat{x}_a = \mathcal{U}^{-1}(u_a) = x_0 - \frac{1}{c} \cdot \log\left(1 - \frac{u_a}{b}\right). \quad (4.9)$$

4.6.4 Choosing Between Conditional Utility Models

These three models are shown for a typical subject in Figure 4-5. Each component modelling technique was used as part of a multilinear utility function to predict the subjects' paired comparison choices. Since there was no difference in the performances of the piecewise linear and the exponential models¹³ the exponential model was chosen for its simplicity and generality.¹⁴




Figure 4-5: A typical subject's conditional utility for time (S3), showing both methods of fitting a utility function to the elicited points: (a) linear interpolation, and (b) an exponential function. A straight-line (c), which represents a hypothetical *risk-neutral* utility function, is shown for comparison.

4.7 Verification of the Independence of Conditional Utilities

Recall that in order to use a multilinear utility function, we require utility independence among the attributes. At several points during the assessment of the conditional utilities, this independence was verified by asking the subject if changes in the levels of any of the other attributes or sets of other attributes would effect the indifference value for the choice between the lottery and the

¹³Both produced errors rates of 7.5%.

¹⁴Although both piecewise and exponential models produced better results than the linear model, there were insufficient data to substantiate the differences.

certain outcome. In every case, the subjects' responses were negative. This result justifies the assessment of conditional utilities and scaling constants required for a multilinear function, instead of the assessment of a full 3-dimensional utility function.

4.8 Determination of the Scaling Constants, $k_x \dots k_{xyz}$

With the component utilities determined, the more difficult job of determining the preference structure *between* attributes remained. When a multilinear utility function is used, this structure is embodied in the scaling constants, $k_x \dots k_{xyz}$. Because of the difficulty of conveying the required ideas to the subjects, who were not experts in decision theory, two methods were used to assess the scaling constants required in Equation 4.1. The first, graphical scaling, proved the easier to explain to the subjects. The second, probabilistic scaling, provided a check, and for several of the subjects it produced better internal consistency.

4.8.1 Graphical Method

In this technique, the subjects were given a scale like that shown on the left-hand side of Figure 4-6 on which was marked the utility of the best possible outcome: $U_{111} = U(\mathbf{x}_{111}) = U_{xyz}(x_1, y_1, z_1) = 1$, corresponding to low flight time, low fuel burn, and no turbulence, and that of the worst possible outcome: $U_{000} = U_{xyz}(x_0, y_0, z_0) = 0$, corresponding to high flight time, low destination fuel, and long exposure to turbulence. They were asked to mark the utilities of the six intermediate corner points ($U_{001}, U_{010}, U_{011}, U_{100}, U_{101}, U_{110}$) in such a way that the position on the scale was proportional to the desirability of the outcome, given the desirabilities of the two end points. The right-hand side of Figure 4-6 shows the scale as returned by a hypothetical subject. The required values were then




Figure 4-6: Graphical assessment of the corner utility points. The subjects were shown the scale in the left hand panel, and were asked to fill it out, producing something like the scale shown in the right hand panel.

measured directly from the scale.

4.8.2 Lottery Method

In this method, the subjects were given the choice between the certain outcome whose utility was to be assessed, say $\mathbf{x}_{100} = (x_1, y_0, z_0)$, and a lottery, $L = \langle \mathbf{x}_{111}, p; \mathbf{x}_{000} \rangle$, and the lottery's probability p was adjusted until the subject was indifferent between the two: $\mathbf{x}_{100} \sim L$. Assuming that preferences

are linear in probability, then the utility of the certainty can be calculated thus:

$$\mathcal{U}(\mathbf{x}_{100}) = \mathcal{U}(L) = p\mathcal{U}(\mathbf{x}_{111}) + (1 - p)\mathcal{U}(\mathbf{x}_{000}) = p. \quad (4.10)$$

This procedure was repeated for each of the six intermediate corner points. This parallels the method used in Section 4.5, except that in this case it was the probability—rather than the bracketing values of the attribute—which was varied.

For each subject the corner points were assessed both ways.¹⁵ Since the assessment procedures were difficult to describe to the subjects, and since there was no *a priori* reason to prefer one method over the other, the better of the two at predicting each subject's paired-comparison choices was used to build the model of his utility function.

4.8.3 Calculation of the Scaling Constants

With the values of the corner points obtained using the better of the above methods, it is possible to calculate the scaling constants by writing Equation 4.1 for each of the assessed corner points:

$$\begin{aligned} \mathcal{U}_{100} &= k_x, \\ \mathcal{U}_{010} &= k_y, \\ \mathcal{U}_{001} &= k_z, \\ \mathcal{U}_{110} &= k_x + k_y + k_{xy}k_xk_y, \\ \mathcal{U}_{101} &= k_x + k_z + k_{xz}k_xk_z, \\ \mathcal{U}_{011} &= k_y + k_z + k_{yz}k_yk_z, \\ \mathcal{U}_{111} &= k_x + k_y + k_z + k_{xy}k_xk_y + k_{xz}k_xk_z + k_{yz}k_yk_z + k_{xyz}k_xk_yk_z. \end{aligned} \quad (4.11)$$

Solving Equations 4.11 in order, with $\mathcal{U}_{111} = 1$, gives the required constants:

$$\begin{aligned} k_x &= \mathcal{U}_{100}, \\ k_y &= \mathcal{U}_{010}, \\ k_z &= \mathcal{U}_{001}, \\ k_{xy} &= \frac{\mathcal{U}_{110} - \mathcal{U}_{100} - \mathcal{U}_{010}}{\mathcal{U}_{100}\mathcal{U}_{010}}, \\ k_{xz} &= \frac{\mathcal{U}_{101} - \mathcal{U}_{100} - \mathcal{U}_{001}}{\mathcal{U}_{100}\mathcal{U}_{001}}, \\ k_{yz} &= \frac{\mathcal{U}_{011} - \mathcal{U}_{010} - \mathcal{U}_{001}}{\mathcal{U}_{010}\mathcal{U}_{001}}, \\ k_{xyz} &= \frac{1 + \mathcal{U}_{100} + \mathcal{U}_{010} + \mathcal{U}_{001} - \mathcal{U}_{110} - \mathcal{U}_{101} - \mathcal{U}_{011}}{\mathcal{U}_{100}\mathcal{U}_{010}\mathcal{U}_{001}}. \end{aligned} \quad (4.12)$$

The coefficients for each subject are presented in Table C.1.

4.9 The Complete Multi-Attribute Utility Function

Combining the scaling constants determined in Equation 4.12 and the component utility functions of Equation 4.7 into Equation 4.1 provides each subject's multilinear utility function of flight time, fuel at destination, and exposure to turbulence. Figure 4-7 shows three views of the four-dimensional

¹⁵One of the subjects was unable to understand the lottery method of assessment, and his component utilities could not be reliably determined.

utility function for a typical subject, each with one of the attributes held constant at its most desirable value. Figure 4-8 highlights the differences between the elicited utility functions for a

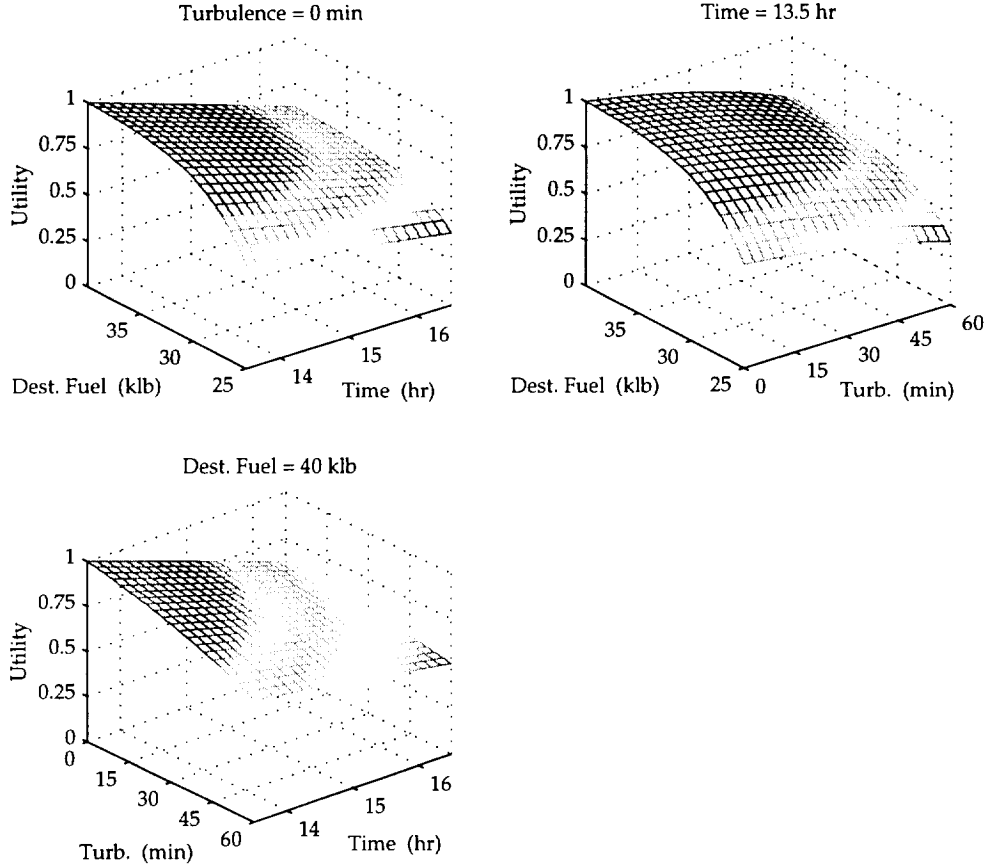


Figure 4-7: Three views of the final multilinear utility function for a typical subject (S_5). For each plot, the omitted variable is held constant at its most desirable value.

typical subject, and the cost function for his aircraft. Iso-preference lines according to the utility-maximization model are shown in the left hand panel, and according to the cost-minimization model in the right hand panel. For the purposes of comparison, exposure to turbulence was held constant at zero for the utility model. Note the marked curvature of the iso-preference lines on the left.

The significance of the difference between the utility and cost models shown in Figure 4-8 is apparent, since the two sets of iso-preference curves are quite different. Even if the cost index were varied, which is equivalent to adjusting the slope of the family of lines in the right hand panel of Figure 4-8, it would not be possible for the cost model to reproduce the preference structure exhibited by the utility model.

4.10 Risk Aversion

As can be seen in Equation 4.7, each conditional utility function is specified completely by two end-points and a single parameter, c . It is this parameter which contains the most valuable infor-




Figure 4-8: Iso-preference curves for two different decision models over flight time and destination fuel for a typical subject (S_5): the multilinear utility function (shown in two dimensions, with turbulence held constant at zero) on the left, and the cost-minimization model on the right. The iso-preference curves are marked in *utils* and thousands of dollars respectively.

mation about a subject's preference structure, but it is not independent of scale.¹⁶ Risk aversion—a measure of the extent to which a DM prefers a certain outcome to a lottery with the same expected consequence—provides a more general measure.

Risk aversion is traditionally defined (see Keeney & Raiffa, 1993, p. 183) as the ratio of the second and first derivatives of utility:¹⁷

$$r_x(x) \equiv \begin{cases} -\frac{\mathcal{U}_x''(x)}{\mathcal{U}_x'(x)} & \text{if } \frac{d\mathcal{U}_x}{dx} > 0, \\ \frac{\mathcal{U}_x''(x)}{\mathcal{U}_x'(x)} & \text{if } \frac{d\mathcal{U}_x}{dx} < 0. \end{cases} \quad (4.13)$$

However, the exponential utility function used in Equation 4.6 is not scaled in the traditional way: it is not necessarily increasing, and does not necessarily start at $x = 0$ and end at $x = 1$. For these reasons it was necessary to redefine risk aversion in the following dimensionless form:¹⁸

$$r_x(x) \equiv -\frac{\mathcal{U}_x''(x)}{\mathcal{U}_x'(x)} \cdot (x_1 - x_0). \quad (4.14)$$

Substituting values of $\mathcal{U}_x'(x)$ and $\mathcal{U}_x''(x)$ from the exponential utility function of Equation 4.7, gives:

$$r_x(x) = c \cdot (x_1 - x_0). \quad (4.15)$$

Risk aversions for each attribute, calculated using Equation 4.15, are shown by subject in Table 4.4

¹⁶Specifically, the scaling of the attribute axis affects the value of c . Since the exponent $c(x - x_0)$ must be dimensionless, the units of c must be the inverse of the units of x . Thus c will depend, among other things, on the units chosen for x —an unsatisfactory situation.

¹⁷For utility functions which are twice continuously differentiable.

¹⁸Note that another solution to this problem would have been to redefine the exponential utility function thus: $\mathcal{U}(x) = (1 - e^{-c'x'})/(1 - e^{-c'})$, where $x' = (x - x_0)/(x_1 - x_0)$. However, this form tends to obscure the scaling of the ordinate values.

and in Figure 4-9.

<i>subj.</i>	r_f	r_t	r_q
1	1.774	1.744	0
2	1.478	7.673	-1.527
3	1.052	2.162	0
4	1.836	4.872	-0.458
5	4.022	0	0.947
means	2.032	3.290	-0.208

Table 4.4: Least-squares risk aversion by attribute and subject. Negative values indicate risk-seeking behavior.




Figure 4-9: Risk aversions by subject and attribute. Note the risk-seeking behavior of some of the pilots in their preference structure for exposure to turbulence.

Of the fifteen mean risk aversions shown in Table 4.5, five are statistically significantly different from zero at the 5% level: r_{f_4} , r_{f_5} , r_{t_4} , r_{q_2} , and r_{q_5} . Also, the only statistically significant risk-seeking behavior observed was subject 2's preference for exposure to turbulence.¹⁹

¹⁹It is interesting that any of the subjects exhibited risk-seeking behavior at all. It might best be explained by the pilot's perception of a premium on exposing passengers to no turbulence—perhaps some passengers consider some turbulence to be quite bad, but more turbulence to be little worse. With such a premium, pilots would be willing to gamble in order to achieve the very desirable outcome of a smooth ride for the entire flight, and would thus appear to be risk-seeking.

<i>subj.</i>	r_f	$p(H_0)$	r_t	$p(H_0)$	r_q	$p(H_0)$
1	1.854	0.062	2.205	0.053	0	1
2	1.912	0.076	12.285	0.052	-1.690	0.017
3	1.206	0.075	3.247	0.084	0	1
4	2.071	0.019	5.690	0.013	-0.528	0.294
5	4.101	0.0003	0	1	0.843	0.049

Table 4.5: Mean risk aversion, and p -values associated with a two-tailed null hypothesis that the mean was zero, by attribute and subject.

4.11 A Comparison of Several Decision Models

During the course of the experiment, several of the subjects indicated that they were choosing between the pairs of flight plans using lexicographic ordering (described in Chapter 2). They were deciding on the value of the attribute which they considered most important, and using other attribute values only as needed to break ties in the first. For each subject, each of the six possible lexicographic decision models²⁰ was used to predict decision behavior, and the best was retained for comparison.

The subjects had also been given a cost index of 100 for their flight. Predictions of their decisions were made on the basis of minimizing the cost calculated using Equation 2.2 with $I_c = 100$.

These three models of human decision making—utility maximization, lexicographic ordering, and cost minimization—were evaluated against the subjects’ answers to the paired comparison questions. Figure 4-10 shows the fraction of errors made by each of the three models of the pilots’ decisions.

The Student t -test was used to make the comparisons between models. The utility model was significantly better at predicting pilot decisions than was the lexicographic model ($p = 0.023$)²¹ and the cost minimization model ($p = 0.00015$). The lexicographic model was just significantly better than the cost model ($p = 0.048$).

From these results, it is clear that the utility model is the best predictor of the subjects’ decisions. This result may be attributed to two of its features: it models the non-linearity of the pilots’ preferences, and it models more of the attributes they take into account when making these decisions.



“Man in sooth is a marvellous, vain, fickle, and unstable subject.” *Michael de Montaigne (1533–1592)*

²⁰There are ${}^nP_n = n!$ lexicographic models in an n -attribute decision space. For the profile selection problem, in which $n = 3$, the six models are all possible permutations of destination fuel, flight time, and turbulence.

²¹All comparisons were made with a null hypothesis that the two means were equal, and an alternate hypothesis that the first mean was lower than the second — all one-tailed tests.




Figure 4-10: A comparison of the predictive performance of the three decision models: utility maximization, lexicographic ordering, and minimization of linear cost. The horizontal gray lines mark the sample means, the vertical gray lines mark the extent of the 95% confidence intervals for the population means; the black boxes mark the 25th, 50th, and 75th percentile points for the samples; and the whiskers mark the 0th and 100th percentile points for the samples. In some cases, several sample points were at the bottom of the range, so that the 0th and 25th percentile points were coincident, hiding some of the whiskers.

Chapter 5

Trajectory Optimization

Maximizing the Objective Function

Putting aside the matter of the selection of an objective function, let us turn our attention to the algorithmic side of trajectory optimization: finding the trajectory or trajectories which maximize the value of an objective function. This chapter begins with a survey of methods which have been applied to the problem, with attention to their advantages and disadvantages. It continues with a discussion of an appropriate representation for the problem, and ends with the development of one particular method for optimizing trajectories: a stochastic directed search, which operates on a discrete representation of the trajectory search space.

5.1 A Survey of Methods

Many methods have been applied to the problem of finding fuel-, time-, and cost-optimal trajectories for aircraft. Several of these techniques are discussed below.

5.1.1 Gradient-Based Techniques

The calculus of variations has been used to optimize the flight profile for a DC-10 (Shaoee & Bryson, 1976). For long-range flights, a profile may reasonably be divided into climb, descent, and cruise segments, and optimization may be performed on the cruise segment using calculus (Lee & Erzberger, 1980; Katz, 1994). While these classical techniques are computationally simple, they are limited in the types of policy they can be used to explore.

A simple method for determining the optimum flight profile is to simply fly at whatever altitude instantaneously provides the lowest time, fuel, or cost given the current winds and aircraft weight. This is the procedure used by the FMS in the Boeing 747-400 (Honeywell, 1994). One limitation of this method, however, is that it does not take into account the additional fuel and time expended during climb and descent.¹

All of these methods produce continuous-climb cruise solutions, which are not compatible with ATC-imposed constraints on altitude or airspeed, and are not necessarily globally optimal. It is also difficult to provide meal service to passengers during a climb or low-air-speed flight² increasing the desirability of having level cruise segments.

¹These quantities are not negligible. According to its Operations Manual (Boeing, 1988), a 760 klb Boeing 747-400 consumes about 7% more fuel and travels about 7% more slowly during climb and descent around 32 kft, than it does in level cruise at 32 kft.

²This is because it can be difficult to pull the food carts uphill at high body angles (Gifford, 1996).

5.1.2 Filtered Cruise-Climb Solutions

This approach to optimization starts with the determination of an optimal cruise-climb profile by any of the aforementioned methods. The resulting profile is often unflyable, since it uses a block of altitudes rather than a single altitude, and since it often involves an unacceptably high number of climbs and descents. To fix these shortcomings, the optimal trajectory is then “filtered”³ to constrain it to use only appropriate altitudes (Lidén, 1992). This technique is guaranteed to produce an




Figure 5-1: An illustration of an optimal cruise-climb trajectory (labeled *raw*), and the result of filtering (labeled *filtered*) by using the *nearest* allowable eastbound flight level.

admissible solution, but while it can—at best—produce an optimal solution, the filtering process is almost certain to degrade the optimality of the cruise climb. Another drawback is that in the presence of winds, this method can produce “spikes”—additional pairs of climbs and descents which don’t contribute significantly to optimality.

5.1.3 A Discrete Representation

Because the required filtering operation results in a loss of optimality, it seems evident that there is a problem with the *continuous* representation of the search-space which is used in the aforementioned methods. It would be more appropriate to use a *discrete* representation of the airspace, one which only uses appropriate altitudes, and therefore need not be filtered to meet ATC requirements or workload constraints.

Given the airspace constraints illustrated in Figure 1-4, it seems logical to represent a profile by a sequence of constant-altitude cruises, with climbs and descents made only as required to move between admissible cruise altitudes. This approach is also appealing from the pilot’s point of view, as pilots are trained to see flight planning in this manner (Jeppeson, 1996b, 1996a).

5.1.4 Optimization by Search

Many standard search techniques can be used to find optimal trajectories in a discrete search space. They can be broadly categorized as exhaustive, informed, or stochastic, as defined below.

Exhaustive Search

Exhaustive search⁴ produces all possible solutions to a search problem, and is therefore *complete*. It is, however, prohibitively time-consuming. Assuming only vertical exploration, no lateral explo-

³It has been pointed out that this operation isn’t filtering in the traditional sense: instead of removing undesired high- or low-frequency components of the trajectory, it adds desired high-frequency components. The word “filtered” is, however, prevalent in the literature.

⁴Named *the British Museum procedure* by the ever-humorous AI community (Winston, 1992).

ration, a flight composed of s segments, each of which can be flown at any of l flight levels will have $n = l^s$ possible unique profiles. Consider the flight shown in Figure 1-2, which is divided into fifteen segments, each of which might be flown at one of perhaps 10 altitudes: there are 10^{15} possible profiles in this search space. If simulation and evaluation of the consequences of a single flight profile took only 0.0001 seconds, exhaustive search for the best solution might still take 3000 years. Clearly exhaustive search is not an option.

Informed Search

There are several solutions to this problem, all of which involve the use of additional information in the search process—hence *informed* search. The best known of these is perhaps *Dynamic Programming* (Bellman & Dreyfus, 1962), in which the principle of *implicit enumeration*⁵ is used to reduce the effort required to span the search space. For a route with s segments and l flight levels, the number of complete profiles to be simulated, n , is of the order of l^2 . For the previous example, with $l = 10$, the number of solutions, n , is 100, which is a significant reduction from the original 10^{15} solutions.

However, the method is not without its drawbacks. Aside from the complexity of keeping track of the progress of the search, constraint information must be added in the part of the algorithm which generates small admissible steps for exploration, making it difficult to change the constraints once the algorithm is written. In addition, dynamic programming is limited in the types of objective functions it can be used to optimize: it can only be used with objective functions which are *separable*.⁶ In the case of profile optimization, this requires that fuel burn on one segment does not affect fuel burn on any other. Since fuel flow depends heavily on aircraft weight (see Figure 1-5) this assumption is not valid. This problem can be circumvented by iterating over final weight (Bellman & Dreyfus, 1962, use this method to find the profile which minimizes time-to-climb for an interceptor), or by augmenting the search space, but these solutions are both complex and computationally intensive.

Combining dynamic programming with a heuristic to direct the exploration towards the most promising part-paths first reduces the search time significantly. In the so-called A^* algorithm, the heuristic is to use an underestimate of the total objective function associated with each part-path to sort the paths before further exploration. This method is very efficient at finding optimal solutions but again suffers from the limitation of one of its components—dynamic programming—which is that it requires a separable objective function. Barrows (1993) used a modified A^* search to find the optimal profile for a *Piper Arrow*. The Arrow's fuel fraction by weight was only about 16%,⁷ so its mass and fuel consumption could reasonably be considered independent of prior fuel burn, making the objective function—a linear combination of fuel and time—separable. By contrast, the Boeing 747-400—a typical long-haul transport aircraft—has a fuel fraction by weight of about 44%,⁸ making this assumption much less tenable. Niiya (1990) used an A^* search to optimize the orbital trajectory of a small spacecraft. He was able to make the same separability assumption because the fraction of propellant by weight was small.

⁵In implicit enumeration, a partial optimization is performed at each stage of the search, progressively reducing the size of the search space.

⁶A separable objective function is one in which the value of the objective function for each stage of the search—in this case, each stage of the flight—is independent of the values of the objective function for every other stage of the search.

⁷The Arrow's usable fuel capacity is 72 gallons or 432 lb, and its maximum gross weight is 2750 lb (Piper Aircraft Corporation, 1978).

⁸The 747-400's fuel capacity is about 383 klb, and its maximum takeoff weight is about 870 klb (Boeing, 1988).

5.1.5 Stochastic Directed Search and the Genetic Algorithm

Because of the complications involved in implementing any of these search algorithms, a simpler approach seemed warranted. The *Genetic Algorithm* or GA (Goldberg, 1989), which is a stochastic directed search loosely based on the evolution of species observed in nature, provides relief from many of the problems described in the previous sections. In a GA, the sexual reproduction of pairs of solutions and the mutation of some individual solutions are used to span the search space stochastically, and selection based on an explicit objective function is used to prune some of the worst individuals to make room for better ones.⁹ The basic operations which make up the simple GA are crossover, mutation, evaluation, and selection.

A good search algorithm is one which uses operators which are complete,¹⁰ non-redundant,¹¹ and informed¹² (Winston, 1992). The operators involved in a genetic algorithm (initialization, mutation, and crossover) are not complete,¹³ and they are redundant, because several individuals in a population can be identical.

However, the GA has the distinct advantages of being relatively simple to implement: it does not rely on any derivative information about the objective function, and it requires simpler trajectory extension algorithms, since most of the constraints are handled in the evaluations, as described in the next section. While undoubtedly viewed as inefficient by traditionalists, it is amenable to parallel computation, and produces workable solutions well before it approaches an optimal one. Traditional search techniques like A^* do not provide such progressive refinement in the solution: one must wait until the end of the algorithm for any solution at all.

Perhaps most importantly, the GA employs a *holistic*¹⁴ approach to evaluation: each trajectory is evaluated as a whole, removing the requirements of separability and monotonicity of the objective function. This allows the technique to be used to optimize a much wider class of objective functions, such as the nonlinear utility functions elicited from pilots in Chapter 4.

Constraint Handling

A significant benefit of the GA is that it can easily handle a wide variety of constraints. In the trajectory optimization problem the constraints can be divided into three groups, each of which is handled in a different way: search-space constraints, aircraft performance constraints, and objective constraints.

Search-Space Constraints ATC imposes altitude restrictions which have the effect of discretizing the altitude components of the search space, and speed restrictions, which remove many of the speed components of the search space. These constraints are best handled in the GA's operators (generation, reproduction, crossover, and mutation), which can be designed to ensure that solutions which do not meet these restrictions are not generated.

⁹There is also an almost-separate body of literature on "genetic programming", which deals with the application of GAs to program generation. However, I do not think that the distinction is a useful one. The two fields are equivalent since any optimal policy may be considered to be a program of action for execution by an operator.

¹⁰Complete operators are capable of spanning the entire search space to find the global optimum.

¹¹Non-redundant operators reach each part of the search space only once.

¹²Informed operators use heuristics to ensure that they concentrate the search on the most promising areas of the search space.

¹³But they are asymptotically complete. That is to say, given sufficient time, the probability that they have not generated all possible admissible solutions approaches zero.

¹⁴I use the word *holistic* in the medical sense, in which it refers to a treatment which acts on the whole, rather than on the parts individually (see the definition in Thomas, 1993), rather than in the lay sense, in which it refers to the philosophy that entities are more than the sum of their parts (see the definition in Sykes, 1982).

Aircraft Performance (Evaluation) Constraints Performance-related constraints cannot be easily handled in the GA's operators, which would have to perform a time-consuming simulation to predict their effects before being able to generate new admissible solutions. They are best handled in the evaluation routine. For example, an aircraft's ceiling increases as fuel is burned off, but is initially a limiting factor, as shown in Figure 5-2. A trajectory which is known to violate an altitude capability constraint can be appropriately modified (or the offending command altitude can be ignored) during the flight simulation.




Figure 5-2: Altitude capability as a function of weight for the Boeing 747-400 with P&W 4056 engines.

Objective Constraints Aircraft are often required to arrive at a certain time, and are always required to land with a specified minimum of fuel. These constraints are not an explicit part of the search space, rather they are constraints on attributes which are used in the objective function, and thus are best dealt with through direct manipulation of the objective function, rather than by the application of more informed GA operators. For example, flight time is determined by very complex combinations of the chromosome parameters, which therefore cannot easily be manipulated individually to adjust schedule adherence.

Several researchers have used GAs to solve guidance problems. Schultz (1991) used a GA to evolve strategies for navigation and collision avoidance for an autonomous submersible craft. van Deventer (1993) used a GA to determine an optimal sequence of aircraft bank angles, effectively finding a horizontal trajectory which would allow the aircraft to pass near certain waypoints while remaining clear of threat areas. Delahaye et al. (1994) used a GA to distribute aircraft between several ATC sectors to reduce sector workload and thus maximize airspace capacity, and Durand et al. (1995) used a GA to provide lateral guidance for conflict resolution between aircraft. However, the GA has not previously been applied to vertical flight-path planning.

5.2 Optimization by Genetic Algorithm

A system was implemented in ANSI-standard C++, following the flow chart shown in Figure 5-3 whose elements are described individually below, to optimize the vertical navigation for a typical route from Los Angeles to Sidney, Australia shown in Figure 1-2.

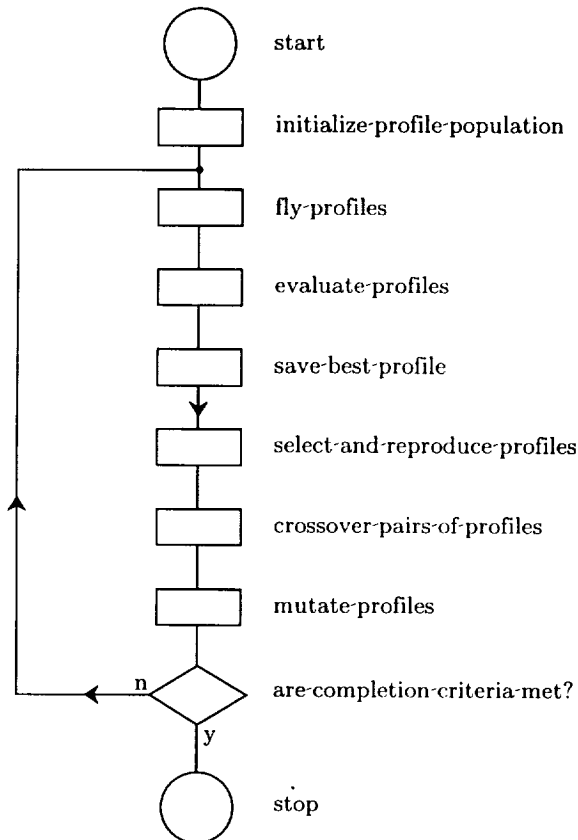


Figure 5-3: A flow chart for the simple genetic algorithm.

Representation & Chromosome The GA operated on the airspace representation discussed in Section 5.1.3: the route was divided horizontally into 15 segments, as shown in Figure 1-2, and distance and course for each of the segments was pre-computed using the algorithms given in Appendix D. The route was also divided vertically into the many allowable flight levels. The GA was designed to find the best sequence of 15 flight levels, using the *command-altitude*¹⁵ representation shown here in Figure 5-4.

Initialization The population, which typically consisted of a few hundred individuals, was seeded with a mix of single-altitude profiles and multiple-altitude profiles to ensure sufficient genetic diversity.

¹⁵The term *command-altitude* refers to the fact that *desired* altitudes are encoded in the chromosome. The simulation limited the actual altitude on a segment according to the altitude capability of the aircraft at the prevailing weight. This representation allowed the GA's operators to be ignorant of the aircraft's performance constraints, making their design much simpler.

31	31	35	35	35	39	43	43	43	43	43	43	43	43	43	43
----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----

Figure 5-4: The chromosome used to represent a typical profile for the fifteen-segment flight. Note that the sequence of command altitudes is read from left to right.

Evaluation by Flight Simulation Each profile was first evaluated by simulating the flight of a typical long-haul aircraft—the Boeing 747-400—over the route at the profile’s specified altitudes. A detailed description of the aircraft performance model used is given in Appendix E.

For segments on which the commanded altitude was above the maximum altitude for the aircraft for the whole segment, the profile was “edited”—the unattainable altitude was reduced to the aircraft’s maximum altitude. This speeded up convergence of the GA significantly: it ensured that the command altitudes were always low enough that any changes that were made to them in the course of the optimization actually affected the profiles as flown.¹⁶

Objective Function Evaluation After a profile was flown, the resultant attributes (e.g. fuel at destination, flight time, and turbulence exposure) were combined using an objective function—such as cost or utility—to produce a single metric of the quality for that profile. This metric was then normalized to produce a fitness index, which was used in the selection and reproduction routines. At this point, the best member of the population was saved for reintroduction after the crossover and mutation operations,¹⁷ which might otherwise have destroyed it.

Selection and Reproduction Individuals were selected for reproduction into the next generation stochastically,¹⁸ and on the basis of their fitness. First, a *scaled fitness*, f_s , was calculated for each individual according to Equation 5.1:

$$f_s = \frac{\mathcal{O} - \mathcal{O}_{min}}{\mathcal{O}_{max} - \mathcal{O}_{min}}, \quad (5.1)$$

where \mathcal{O}_{min} and \mathcal{O}_{max} are respectively the lowest and highest values of the objective function found in the population.¹⁹ Next, *roulette-wheel selection* (as described in Goldberg, 1989) was used to pick individuals for reproduction into the next generation. In this way, more fit individuals were more likely to be reproduced, and were therefore more and more frequently represented in subsequent generations.

Crossover Adjacent pairs of profile chromosomes²⁰ were subjected to crossover with a probability of p_c . If a pair were to be crossed over, a crossover point, P , was chosen along the length of the chromosome using a uniform random distribution, as shown in Figure 5-5. The crossover involved the creation of two new chromosomes, one with the head (everything up to location P) from the

¹⁶For example, if the command altitude was 10,000 feet above the aircraft’s maximum altitude, then any change of a few flight levels in the command altitude would result in no change at all in the altitude at which the aircraft flew, which would still be its maximum altitude.

¹⁷This strategy is called *elitism*.

¹⁸To avoid the problems often associated with system-supplied random-number generators (e.g. sequential correlation) I implemented a version of the “shuffling” random-number generator attributed to Bays and Durham, in Press et al. (1988).

¹⁹This scaling ensures that the selection process can work both at the beginning of the algorithm, when there is a large difference between the best and worst trajectories in the population, and at the end as the population converges, when this difference is relatively small.

²⁰Adjacent in the array structure used to hold the population in the program.

first parent and the tail (everything from location P on) from the second, the other with the head from the second and the tail from the first.




Figure 5-5: Single-point crossover of two chromosomes, A and B , to produce two offspring, A' and B' .

Mutation Two types of mutation were employed to prevent the premature loss of potentially good solutions from the population.²¹ The traditional *point* mutation, shown in Figure 5-6, in which a single gene (altitude) is changed at random, was used with probability p_{m_p} on each element of each chromosome. In addition, a *block* mutation, shown in Figure 5-7, was used on each chromosome with probability p_{m_b} . In this second type of mutation, two sites— P_1 and P_2 —were chosen in the chromosome at random (again using a uniform probability distribution for site selection) and all altitudes between them were raised or lowered a small random number of flight levels. The block mutation was added because point mutations are fairly ineffective at producing beneficial changes in a profile. This is because the addition of—say—a closely spaced climb and descent is unlikely to improve the fuel consumption and flight time of a profile.




Figure 5-6: Point mutation of a chromosome, C , to produce a new chromosome, C' .




Figure 5-7: Block mutation of a chromosome, D , to produce a new chromosome, D' .

²¹It is interesting to note that most viruses evolve *only* by mutation—they cannot exchange genetic material by crossover as many organisms do during sexual reproduction. This severely limits their rate of adaptation.

Clean-up by “Hill-Climbing” The directed nature of the GA’s search comes from the distribution of individuals in the population, rather than from any intelligent operations performed on individuals. For this reason, the GA is unlikely to remove small imperfections in a profile, except by chance. Some “post-processing” was therefore required to clean up any such imperfections in the best profile found by the GA. This was achieved by successive unidimensional (single-segment) optimizations: for each segment the altitude was cycled through all allowable flight levels, and that which produced the best trajectory was kept. This process was performed over the whole flight, and repeated until convergence, which was usually after between one and three cycles. These cycles were quite time-consuming, since each required *ls* complete flight simulations.²²

5.3 Results

After the GA’s adjustable parameters had been tuned²³ to the values shown in Table 5.1, the GA was run to produce several pareto-optimal trajectories. The GA was used to produce the minimum

<i>parameter</i>	<i>value</i>
population size, N	150
max no. generations, n_{max}	300
block mutation prob., p_{mb}	0.05
point mutation prob., p_{mp}	0.02
crossover prob., p_c	0.5

Table 5.1: Parameters used in the GA to produce the results shown in the figures below.

fuel profile shown in Figure 5-8, the minimum time profile shown in Figure 5-9, the minimum cost profile shown in Figure 5-10, the minimum cost profile with no winds shown in Figure 5-11, and the maximum utility profile (using the multilinear utility function of Subject 5) shown in Figure 5-12. For an explanation of the symbology used in the figures, see the caption to Figure 5-8.

5.3.1 Excessive Climbing and Descending

Most optimization schemes produce an excessive number of climbs and descent in their solutions. These additional climbs and descents are undesirable because they add to pilot workload, they add to the number of cycles the engines are subjected to, and they make life in the cabin less pleasant. Most methods deal with this problem by adding a constraint, either on the number of climbs and descent, or on their proximity (Lidén, 1992; Barrows, 1993). This is a poor strategy, since such constraints do not take into account the costs associated with each climb and descent. A better way of handling the problem is to add a penalty for climbing or descending to the objective function:²⁴

$$\mathcal{O} = f + \frac{I_c}{10} \cdot t + \frac{1}{2} \cdot (n_c + n_d - 2), \quad (5.2)$$

where I_c is the airline’s cost index, and n_c and n_d are the number of climbs and descents respectively in the profile.²⁵ Minimizing Equation 5.2 with $I_c = 100$ produces the optimal profile shown in

²²This method of optimization alone is not sufficient to find optimal trajectories: it can converge on a local optimum, as was verified experimentally during the optimizations described in the next section.

²³The parameters were tuned “by eye”. See the proposal for a better method in Section 7.2.3.

²⁴Note that *performance* penalties for climb and descent are already built into the flight simulation routine.

²⁵Every flight is—hopefully—composed of one climb and one descent, so 2 is subtracted from the total number of climbs and descents, giving the number of excess climbs and descents in the flight.




Figure 5-8: The minimum-fuel ($\mathcal{O} = f$) trajectory produced by the GA. Note the winds—shown in gray—each acting over the third of the route up to its vertical gray line. Note also the difference between the commanded profile in gray, and the profile actually flown in black.

Figure 5-13, which should be compared with the standard minimum-cost profile shown in Figure 5-10. The best possible way of handling the climb-and-descent problem would be to build a multi-attribute utility function which modelled the pilots preferences over the number (and perhaps length and spacing) of climbs and descents in addition to the attributes considered in Chapter 4, and then find the profile which maximizes that utility function. Elicitation of such a utility function was beyond the scope of this research.

5.3.2 Schedule Adherence

The GA can also be used to help an aircraft meet a schedule constraint. Figure 5-14 shows the profile devised to minimize the following objective function:

$$\mathcal{O} = f + k \cdot \text{abs}(t - 16), \quad (5.3)$$

with $k = 30$. Minimizing this objective function almost produces a minimum-fuel profile, with the $k \cdot \text{abs}(t - 16)$ term driving the schedule error towards zero. Another objective function which would provide a level of schedule adherence is:

$$\mathcal{O} = \begin{cases} f & \text{if } \text{abs}(t - 16) \leq \tau, \\ f + k \cdot \text{abs}(t - 16) & \text{if } \text{abs}(t - 16) > \tau. \end{cases} \quad (5.4)$$

where τ is some small time interval within which schedule adherence is not considered a factor.




Figure 5-9: The minimum-time ($\mathcal{O} = t$) trajectory produced by the GA.

5.3.3 Airspace Design

The optimization system was also used to explore the consequences of the discretization of the vertical structure of the airspace. Such discretization is used to provide vertical separation between aircraft, because altitude has traditionally been easier to measure and prescribe precisely than has location, and because aircraft have traditionally flown along routes defined by expensive ground-based navigation aids.²⁶ Due to concerns about inaccuracy in altimetry, the separation between opposite-direction aircraft increases from 1000 ft below 29 kft, to 2000 ft above 29 kft. As altimetry is improved, this vertical separation can be reduced.

To examine the effects of such a relaxation of separation standards, the GA was run with each of the following airspace structures: (1) with ICAO altitudes (2000 and 4000 foot separations between same-direction traffic below and above 29 kft respectively) as shown in Figure 5-16, (2) with 2000 separations at all altitudes as shown in Figure 5-15, and (3) with 1000 foot separations at all altitudes as shown in Figure 5-10.²⁷ In an attempt to remove any bias, three different wind conditions were used: (a) the winds shown in Figure 5-10, (b) the same winds with their direction reversed, and (c) no wind (i.e. calm). The results are shown in Table 5.2. The demonstrated increase of about one

<i>separation</i>	<i>figure</i>	<i>% increase</i>
1000 ft	5-10	0
2000 ft	5-15	0.13
ICAO	5-16	0.43

Table 5.2: The increase in cost associated with several different vertical structures of the airspace. Note that percentage changes are based on the 1000-ft separation cost.

²⁶With modern navigation techniques such as GPS, and separation techniques such as TCAS, such vertical separation might be replaced by horizontal separation, allowing aircraft to better optimize their flight profiles.

²⁷Note that 1000-ft separation is as close as the GA's representation can get to simulating continuous profiles.




Figure 5-10: The minimum-cost trajectory ($\mathcal{O} = f + \frac{I_c}{10} t$, $I_c = 100$) produced by the GA.

half of a percent for the ICAO separations represents a substantial additional cost.




Figure 5-11: The minimum-cost trajectory ($\mathcal{O} = f + \frac{I_c}{10} t$, $I_c = 100$) with no winds produced by the GA.




Figure 5-12: The maximum-utility (for S_5) trajectory produced by the GA.




Figure 5-13: The minimum-cost trajectory with limited climbs and descents ($\mathcal{O} = f + \frac{I_c}{10} t + 0.5 (n_c + n_d - 2)$, $I_c = 100$) produced by the GA.




Figure 5-14: The minimum fuel trajectory with schedule adherence ($\mathcal{O} = f + 30 \text{ abs}[t - 16]$) produced by the GA.




Figure 5-15: The minimum-cost ($\mathcal{O} = f + \frac{I_c}{10} t$, $I_c = 100$.) trajectory using 2000 foot separations.




Figure 5-16: The minimum-cost ($\mathcal{O} = f + \frac{I_c}{10} t$, $I_c = 100$.) trajectory using ICAO separations.

Chapter 6

A Decision Support System *for Trajectory Optimization*

6.1 The Current Situation

Subject to Air Traffic Control (ATC) separation constraints, airline dispatchers and pilots have some latitude in selecting cruise profiles (Wagenmakers, 1991). Extensive flight-planning tools are available to dispatchers for pre-flight use to aid in the performance of this task. However, once in flight, pilots have only very limited tools available to help them evaluate the attributes of a proposed trajectory. While the FMSS in some aircraft (e.g. the Boeing 757 and 767) have the capability to evaluate the consequences of flying a single proposed altitude other than that actually being flown, other FMSS do not even have this limited capability (Midkiff, 1996).

For example, there is a wind-altitude trade table in the B-747-400's Operations Manual (Boeing, 1988) which presents the break-even head-winds for different altitudes, but it does not account for the time and fuel consumed in climbing to and descending from the proposed altitude.

In the exploratory survey described in Chapter 3, the pilots were asked what additional features they would like in their FMSS to help them with the tasks of vertical navigation and altitude selection. Considering the data of Table 3.6, however, it seems that pilots desire much more information than is currently available. Such information can also be very valuable operationally: the NASA Cockpit Weather Information (CWIN) system (Nolan-Proxmire et al., 1996), which provides pilots with real-time weather information to help in their decision making, has produced fuel and distance reductions of 5% in an MD-11.

6.2 Design Considerations

Automation or Decision Aiding While it is tempting to automate the problem of trajectory optimization, perhaps using a composite utility function and some capable optimization algorithm, such a system would effectively take the pilot out of the decision-making loop, leaving him with no opportunity to consider unmodelled factors in choosing a flight path. There is also evidence to suggest that increasing the level of automation for a task in which the pilot is capable of performing at some reasonable level does not necessarily improve task performance (Wiener et al., 1991). These factors help make the case for decision-aiding, rather than automation, for trajectory planning.

Visual or Textual Display The literature abounds with examples of the fact that visual displays confer on the human operator a far greater ability to assimilate information than do textual ones

(e.g. Chandra, 1989, who studied the presentation of flight clearances). Much work has also been focused on the type of visual display used in the control of vehicles. In a study of displays for the control of high-speed trains, Askey (1995) found that advisory displays provided the lowest overall workload, and the best performance.¹

Representation Research has increasingly focused on providing the human operator with situation displays and decision aids which make use of a concrete representation of the control or decision problem: the *search space* of Figure 2-1. This may be appropriate in many cases. However, in the case of a decision support system for a complex multiattribute problem, particularly one in which the decision and search spaces are so different, it should be more appropriate to present information in a more abstract representation: the *decision space* of Figure 2-1.

6.3 The Parallel-Axis Display

Consider the five trajectories presented in Table 6.1. These typical trajectories might be pareto-optimal trajectories generated by an optimization algorithm, from which the pilot is expected to choose the most desirable. In this tabulated form it is quite difficult to pick a suitable alternative. These trajectories could be displayed on a conventional x - y plot, as shown in Figure 6-1, but

<i>Dest. Fuel</i> (klb)	<i>Flight Time</i> (hr)	<i>Turbulence</i> (min)
29.208	15.823	10
25.146	15.288	13
29.665	15.708	20
26.500	15.750	9
28.000	15.475	0

Table 6.1: Tabular data on destination fuel, flight time, and exposure to turbulence, for five candidate trajectories.

information about the third axis must then be shown either textually, or using some awkward visual representation of the third dimension. An alternative form of display is the *parallel-axis plot* (de Neufville, 1990; Keeney & Raiffa, 1993; Bassett, 1995), in which each point in n -space is transformed into a segmented line joining the n component values on each axis. The five trajectories of Table 6.1 and Figure 6-1 are shown in the parallel-axis plot of Figure 6-2.²

Such a parallel-axis plot can be used to display points in an arbitrarily large decision space, but it does not tell the decision maker much about the relative importance of the attributes. However, armed with some information about this relative importance—determined either from a knowledge of the user's or the designer's utility function,³ or from a cost function such as that used by the airlines—it is a simple matter to scale the axes accordingly. Figure 6-3 shows the same trajectory data in another parallel-axis plot, in this case with the axes scaled by Subject 5's multilinear utility function constants k_f , k_t , and k_q .

¹She also found that these same high-performing displays resulted—undesirably—in the highest “head-down” time. During head-down time, the human controller cannot pay attention to factors in his environment which are not presented on the displays.

²This display was used to show the subjects in the utility assessment experiment of Chapter 4 several of the pairs of alternative trajectories they were presented with textually. Every one of them preferred this graphical display to the textual.

³A composite utility model, built using constants determined by averaging those from a number of pilots might also be used scale the axes.




Figure 6-1: A two-dimensional plot of the three-attribute options available to the decision maker.

In comparing Figures 6-2 and 6-3, it becomes apparent that flight time is a less important outcome than the other two, and that the differences between the candidate trajectories are diminished accordingly. A further improvement to the display is possible if we use the risk aversions elicited from the subject to distort internal measures along the axes. Figure 6-4 shows such a parallel-axis plot. It should be clear from this presentation that—according to Subject 5—the trajectories are all fairly good when considering flight time, but vary markedly when considering destination fuel.

Finally, a further axis could be provided in which to display the residual utility from Equation 4.1. This residual utility represents the terms $k_{ft} k_f k_t \mathcal{U}_f(f) \mathcal{U}_t(t) \dots k_{ftq} k_f k_t k_q \mathcal{U}_f(f) \mathcal{U}_t(t) \mathcal{U}_q(q)$, which are not already incorporated into the display of Figure 6-4.

From any of the parallel-axis plots it should be easy to judge whether or not one candidate trajectory is dominated⁴ by another. If the lines representing two trajectories do not cross, then the lower trajectory is dominated by the upper trajectory.

◇ ◇ ◇

⁴One solution is said to dominate another if it is better in every individual attribute.




Figure 6-2: A parallel-axis plot, which is capable of showing the levels of attainment of many attributes.




Figure 6-3: A second prototypic parallel-axis plot, in which the axes have been sized according to the relative importance of each attribute.




Figure 6-4: A third prototypic parallel-axis utility plot, in which the axes have been sized according to the relative importance of each attribute, and distorted to show the nonlinearity of the subject's preference over each attribute.

Chapter 7

Conclusions, Contributions, and Suggestions for Further Research

7.1 Conclusions & Contributions

7.1.1 Decision Theory

Value-Based Modelling Because of concerns about safety, most decisions made in the cockpit are made procedurally (i.e. by following airline- or FAA-mandated procedures), and thus are not amenable to value-based modelling. Many decisions which might be modelled using value involve safety: they involve outcomes which are so unlikely, and so extreme in nature, as to be incomprehensible to the human decision maker. These features render them poor candidates for value-based modelling. Vertical profile selection, however, has been shown to be an area in which value-based modelling can be applied successfully. This is the first known use of value-based decision modelling for this application.

Utility Theory Utility theory provides a better model of pilot decision making for the specific multi-attribute profile selection problem examined herein than do the other behavioral and prescriptive models—most notably the airlines' linear cost model—considered. This is because (a) the resources in the decision problem are constrained, making preferences markedly nonlinear; and (b) the decision makers are considering more attributes of the problem than some of the simpler models account for.

Risk Aversion Risk aversion provides a simple measure of a decision maker's preferences over an attribute. A modified measure of risk aversion—one which is insensitive to the scaling of the attributes—was developed and was used to examine the pilots' preferences. Overall, while there were significant differences between pilots, they were found to be moderately risk averse for destination fuel, inconsistently risk averse for flight time, and essentially risk neutral for exposure to turbulence.

Pilots and Dispatchers Pilots and airline dispatchers have been shown to have similar preferences when prioritizing safety, ride quality, destination fuel, schedule adherence, fuel burn, and flight time. Not surprisingly, pilots were observed to place a slightly higher emphasis on destination fuel, with dispatchers placing more emphasis on schedule adherence.

A General Framework for Decision Modelling A general framework for decision modelling has emerged. It consists of three operations: (1) the selection of one or more candidate points in the

search space, (2) the evaluation of some or all of the selected points in the search space, to produce candidate points in a *decision space*, and (3) the comparison of some or all of the points in the decision space using a multi-attribute value or utility function to produce points on a *value* or *utility* scale. This framework is seen as encompassing many models of decision making, from traditional value-based models such as utility theory, to very behavioral models such as recognition-primed decision making.

7.1.2 Optimization

A genetic algorithm (GA) has been developed which can be used to optimize the vertical component of an aircraft's trajectory. It can handle a wider class of objective functions than can other optimization methods since it does not require the same assumptions about the objective function (e.g. separability and monotonicity). The GA also reduces the difficulty of including constraints, which can be handled in any of three different ways: as search-space constraints, as performance constraints, or as decision-space constraints. The GA's ease of implementation make it a powerful tool for exploration of new airspace designs, air traffic management policies, and decision-aiding techniques. This is believed to be the first application of a GA to the problem of vertical profile optimization.

7.1.3 Airspace and Operations

The cost penalty associated with the present structure of the airspace (i.e. vertical separation in 4000-foot intervals above 29,000 feet) is significant, at about 0.5% excess cost compared to continuous (non-discretized) airspace (as conservatively approximated in this research using 1000-foot intervals). This argues for further relaxation of vertical separation standards, perhaps in favor of the provision of separation laterally.

7.2 Suggestions for Further Research

7.2.1 Decision Theory

While the constantly risk-averse exponential model described in Section 4.6.3 for the conditional utility functions fit the elicited preference data well, there was some evidence of increasing risk aversion in the data. It would be illuminating to apply models which take this behavior into account, and to test them against the models used herein.

7.2.2 Decision Aiding

It would be valuable to build a prototype of the decision aid described in Chapter 6, and evaluate the differences between decision-making performance (as judged against the quasi-prescriptive criterion of utility maximization) with a search-space display and a decision-space display of the profile information. Such a display has the potential to simplify the presentation of multiple candidate solutions to complex multi-attribute decision problems like this trajectory selection problem.

7.2.3 Optimization

Including Airspeed in the Representation Further research into the airspace policy issues touched upon in this research would require an increase in the capability of the optimization software. This could be achieved by including airspeed in the GA's representation (as shown in the proposed

chromosome in Figure 7-1), the GA’s operators, and the aircraft performance model. It might also be beneficial to include variable segment length in the representation.

<i>Altitude</i>	31	31	35	35	35	39	43	43	43	43	43	43	43	43	43
<i>Mach No.</i>	.84	.84	.86	.86	.86	.86	.85	.85	.85	.85	.83	.83	.83	.83	.83

Figure 7-1: The augmented chromosome proposed to represent both altitude and speed for a profile. Note that the speed and altitude parts of the chromosome would be manipulated separately—they are only shown together for clarity.

In-Flight Reoptimization There are two ways to use an optimization scheme for flight planning: it can be used only before the flight (which is then flown “open-loop”) to figure an optimal trajectory based on forecast winds and temperatures, or it can be performed in flight (in addition to before the flight), taking current winds and temperatures into account. This second strategy might be called *inflight reoptimization*.¹ A study might be designed—using actual wind and temperature data—to discover what the long-run average benefits from inflight reoptimization would be.

A Meta-GA for Parameter Tuning A GA may have many parameters, including real-valued parameters like crossover probability and mutation probability, integer parameters like population size and the number of sites used in the crossover operation, and binary parameters like those used to control elitism and profile editing (as done herein to meet performance constraints). These adjustable parameters themselves constitute a substantial search space in which the implementer of a GA must find an appropriate point. This is currently done by tuning the parameters heuristically. Instead, these parameters could be combined into a *meta*-chromosome—one which would code for the type of GA used for a particular optimization problem. A *meta*-GA could then be run to find the best combination of parameters for that problem. This approach would effectively change the designer’s focus from specifying parameters of the problem-solving GA to specifying the objective function used in the parameter-finding GA. Such an objective function might include several attributes of the problem-solving GA it is trying to evolve, such as speed, computational requirements, and the quality of the final solution.

◇ ◇ ◇

¹These two kinds of optimization are sometimes called *open-loop* optimization and *closed-loop* optimization respectively (Bertsekas, 1995).

Appendix A

Preliminary Decision Experiment

This preliminary experiment, which is described in more detail in Patrick (1993), was modelled on one performed by Yntema and Klem (1965). It was designed to examine the possibility of applying utility theory to cockpit decision making. Briefly, the experiment consisted of asking pilots to choose between pairs of alternate airports on the basis of several relevant attributes, and then comparing their decisions with those made according to several decision models, including a utility model.

A.1 Background

Experienced pilots¹ were instructed to imagine themselves on a long cross-country flight in Instrument Meteorological Conditions (IMC) in an aircraft of their choosing, when the destination airport became unusable (for an unspecified but non-weather-related reason). The conditions for the flight were specified as completely as possible: there were neither thunderstorms nor icing conditions, there was unlimited visibility below the an extensive overcast, and these conditions were not expected to deteriorate. It was well before sunset. Fuel was limited by the fact that the questions were asked as though the subject were 1.5 hours into a flight which had been started with full tanks. The alternate airports were all equally inconvenient for reaching the ultimate destination, as shown in Figure A-1.

A.2 Method

A.2.1 Paired Comparisons

Each subject was instructed to choose between three dozen pairs of alternates based on the distance to, and the ceiling at, each alternate. For example, subjects were asked to choose between the following pair of alternate airports: one at 200 n.mi. with a 1000-ft ceiling, and another at 50 n.mi. with a 500-ft ceiling. The subjects were presented with these choices textually, thus:

[200 n.mi., 1000 feet] or [50 n.mi., 500 feet].

The subjects were also asked to indicate the subjective degree of difference between the alternates on a continuous numerical scale of 0 to 7.

¹Mean flight experience for the group was 1345 hours.




Figure A-1: A graphical depiction of a typical situation presented to the subjects. Note that the alternates were always located on the dashed circle, so as to be equally inconvenient for subsequent travel to the ultimate destination.

A.2.2 Utility Assessment

A two-attribute multilinear utility function was assessed for each subject independently of the paired comparisons. The two-part procedure used involved (i) determination of the two conditional utility functions—one over ceiling, the other over distance—and (ii) determination of the scaling constants required to stitch the two unidimensional functions together into a two-dimensional function.

Determination of Component Utilities

Subjects were given two scales, one for distance and one for ceiling, as shown in the left-hand and center panels of Figure A-2, and instructed to mark three to five intermediate points on each scale in such a way that, for example, the position of a mark on the ceiling scale was proportional to the safety of an airport with that ceiling, given a constant value of distance. Anchor points for the best and worst values of each attribute were marked on the scales. Values such as $U_c(500)$ and $U_c(2000)$ were measured directly from the central scale in Figure A-2, and linear interpolation was used to build a piecewise-continuous utility function, $U_c(c)$, from them.

In order that the utility of the distance-related variable be increasing, distance was subtracted from aircraft range to provide range remaining. Also, to account for the substantial differences between different subjects' aircraft, range remaining was then non-dimensionalized with aircraft range, to produce normalized range remaining:

$$r = \frac{R_{max} - d}{R_{max}}. \quad (\text{A.1})$$

Values of utility were determined from the left-hand scale in Figure A-2, using Equation A.1. A piecewise-continuous utility function, $\mathcal{U}_r(r)$, was built from these elicited points using linear interpolation.

Determination of Scaling Constants

Subjects were then given a scale with worst and best anchor points: [worst ceiling, worst distance] (with utility, $\mathcal{U}_{00} = 0$) and [best ceiling, best distance] (with utility, $\mathcal{U}_{11} = 1$) respectively already




Figure A-2: The scales used in the utility assessment, as returned by a typical subject.

marked, and instructed to mark on it the two intermediate points: [best ceiling, worst distance] (with utility, \mathcal{U}_{10}), and [worst ceiling, best distance] (with utility, \mathcal{U}_{01}) in the same manner as before. This scale is shown in the right-hand panel of Figure A-2, and is labelled “Corner Points”. The values of these corner-point utilities were measured directly from this scale.

Construction of a Multilinear Utility Function

These graphically determined utility “corner” points were used to build a multilinear utility function (Keeney & Raiffa, 1993) of the form:

$$\mathcal{U}_{cr}(c, r) = k_c \mathcal{U}_c(c) + k_r \mathcal{U}_r(r) + k_{cr} k_c k_r \mathcal{U}_c(c) \mathcal{U}_r(r). \quad (\text{A.2})$$

Writing Equation A.2 once for each of the non-zero “corner” points gives these three equations:

$$\mathcal{U}_{10} = k_c \mathcal{U}_c(c_1), \quad (\text{A.3})$$

$$\mathcal{U}_{01} = k_r \mathcal{U}_r(r_1), \quad (\text{A.4})$$

$$\mathcal{U}_{11} = k_c \mathcal{U}_c(c_1) + k_r \mathcal{U}_r(r_1) + k_{cr} k_c k_r \mathcal{U}_c(c_1) \mathcal{U}_r(r_1). \quad (\text{A.5})$$

Since $\mathcal{U}_c(c_1) = 1$ and $\mathcal{U}_r(r_1) = 1$ by definition, the three constants k_c , k_r , and k_{cr} are easily determined:

$$k_c = \mathcal{U}_{10}, \quad (\text{A.6})$$

$$k_r = \mathcal{U}_{01}, \quad (\text{A.7})$$

$$k_{cr} = \frac{1 - k_c - k_r}{k_c k_r} = \frac{1 - \mathcal{U}_{10} - \mathcal{U}_{01}}{\mathcal{U}_{10} \mathcal{U}_{01}}. \quad (\text{A.8})$$

For the typical subject whose scales are shown in Figure A-2 (for whom $R_{max} = 350$ nm),

$k_c = 0.35$, $k_r = 0.67$, and $k_{cr} = -0.0682$. Substituting these values into Equation A.2 gives:²

$$\mathcal{U}_{cr}(c, r) = 0.349\mathcal{U}_c(c) + 0.667\mathcal{U}_r(r) - 0.0159\mathcal{U}_c(c)\mathcal{U}_r(r). \quad (\text{A.9})$$

Each subject's utility function was then used to predict his or her choices from the paired comparisons. The results are presented in Section A.3.1 below.

A.2.3 Discriminant-Function Models

The subjects might have been using linear combinations of the alternatives' attributes to discriminate between alternatives in a decision, rather than non-linear functions of the attributes, (e.g. a utility model). To examine this possibility, several linear discriminant-function models were tested on the data. Each discriminant-function model attempts to find a hyperplane which will separate the two classes of decision points (those in which the first alternative was chosen, and those in which the second alternative was chosen).




Figure A-3: A perceptron, or single-element neural network, which was trained to find a discriminant function (hyperplane) to fit each subject's paired-comparison data.

A perceptron algorithm was written in LISP, and run with each of the subjects' decisions as training sets. The perceptron is a very simple neural network. As is shown in Figure A-3, it consists of a set of inputs, x_i , a corresponding set of weights, w_i , a summing junction, and a threshold. Once trained, its output, y , is 0 for the greatest possible number of points on one side of the discriminant hyperplane, and 1 for the greatest possible number of points on the other side of the hyperplane.

The perceptron was run with four inputs (r_1 , c_1 , r_2 , and c_2) to find the best discriminant hyperplane in four dimensions (called the 4-D case). To examine the possibility that the subjects might have been using fewer than the full 4 variables available to them in each decision, the algorithm was also run using only 2 variables for each subject. The variables used for this 2-D case were the difference in normalized ranges remaining:

$$r' = r_1 - r_2, \quad (\text{A.10})$$

²It is interesting to note that the utility function for this subject (S_4) is very nearly *additive*, since $k_c k_r k_{cr} \approx 0$. The assumption of additivity would simplify the utility function to $\mathcal{U}_{cr}(c, r) = 0.349\mathcal{U}_c(c) + 0.667\mathcal{U}_r(r)$. Strictly, the two constants should sum to 1 under the additivity assumption. Thus it would be possible to test statistically for additivity using the graphical assessment method by repeating the "corner-point" assessment l times, and constructing a variable $a_i = k_{ci} + k_{ri}$, $i = 1 \dots l$. The null hypothesis, $H_0: a = 1$, could then be examined using a t test.

and the difference in ceilings:

$$c' = c_1 - c_2. \quad (\text{A.11})$$

The results are presented in Section A.3.1 below.

A.3 Results

A.3.1 Error Rates

Disagreement was measured using Δ , the fractional disagreement between choices the subjects made, and the choices made with their utility functions:

$$\Delta = \frac{1}{n} \sum_{i=1}^n \delta_i, \quad (\text{A.12})$$

where n is the total number of comparisons, and δ_i takes the value 1 if the i^{th} decision produced a disagreement between the models, and 0 if the models were in agreement.

<i>Model</i>	$\overline{\Delta}$
Utility maximization	0.16
4-D Perceptron	0.08
2-D Perceptron	0.13

Table A.1: Mean fractional disagreement, $\overline{\Delta}$, (averaged over all subjects) between the decisions made by the subjects and those made according to the three models.

As can be seen in Table A.1, the utility maximization model does not perform better than either of the discriminant-function models, in spite of being built using a larger number of adjustable parameters. It should be noted that the error rates observed using the utility model on the profile selection problem of Chapter 4 were only half as large as those observed here.

A.3.2 Subjective Difference

The subjective degree-of-difference value recorded in the utility assessment experiment should be a measure of the ease of the decision between the two alternatives. This measure was plotted against difference in utility in the case of the utility model, and against distance from the decision plane for the other models. The correlation coefficients for a linear relationship between these pairs of difference measures are shown in Table A.2.

<i>Model</i>	Mean R^2
Difference in utility	5.6
4-D Perceptron	6.5
2-D Perceptron	0.9

Table A.2: Correlation coefficients for linear relationships between objective distance measures and subjective difference between the alternatives for the three models.

Surprisingly, there is little correlation between any of these measures. Efforts to correlate the subjective difference with the difference in distances alone produced similar results.

A.4 Conclusions

These results indicated that the area in which we were applying value based modelling might not be an appropriate one, and led to the conclusions mentioned in the opening chapters—that value based decision models should not be applied to safety related decisions.

◇ ◇ ◇

Appendix B

Survey Data

B.1 Exploratory Survey

This first questionnaire was designed to discover which elements of the outcome of a flight pilots considered important, and what additional features they would like in their flight management systems. This information forms a basis for the construction of a behavioral decision model in Chapter 4, and for the establishment of requirements for a decision aid in Chapter 6.

<i>Attribute</i>	<i>frequency</i>
Ride quality	27
Weather...	21
unspecified	8
destination	7
enroute	4
departure	2
Fuel burn	20
Time...	20
of flight	16
unspecified	3
in the hold	1
Safety	17
Schedule adherence	12
Fuel at destination	10
Economy/Efficiency	7
Fuel	6
Alternates	6
Marketing factors	1
Crew fatigue	1
Aircraft maintenance condition	1

Table B.1: The attributes of the outcome mentioned by the 32 subjects, and the frequencies with which they were mentioned.

B.1.1 FMS Features Desired by Pilots

Listed below are the FMS features desired by each pilot in the first survey. Also shown are the principal aircraft the pilots were flying at the time of the survey.

s	Aircraft	Comments
16	A300	In the A300 we have a map display. I would like to see more information presented. Information like boundaries, FIR's, minimum descent altitude, height of terrain, warning areas, etc.
22	A300	An easy reference page for time-to-altitude predictions, that requires less "head-in-cockpit" time. Likewise wind/altitude change calculations take too long when choices of altitude are requested by ATC on over-water routes. Most altitudes on over-water routes are difficult to change. SATCOM would be very useful for direct communication with ATC for altitude requests and weather deviations. Note: FMS capacity needs to be enlarged to accommodate more significant terrain features.
30	A300	FMS that could send and receive information from other aircraft as well as ground stations—real time information. Temp/wind/g-forces (ride)/change in temperature over a period of time. Same for wind. Prediction by FMS from all inputs for down line ride.
17	A300-600	(1) New features in the FMS would help if they give g-force protection information. (2) Over water, an ability to "downlink-uplink" information to the "wind/prog" or "winds aloft" portion of the FMS when not in VHF contact with ground stations (i.e. HF ACARS or SATCOM ACARS uplinks).
23	A300-605R	Should more accurately depict the real capability of the aircraft.
15	B-727	If it could figure the time remaining until you could climb to your optimum—or a new—altitude.
7	B-757	(1) Continuous and separate display of: max. altitude based on performance (we can get this through a trick); max. altitude based on 1.3 g buffet margin (we have to look this up); optimum altitude (we have this now). (2) A planning mode that would allow us to program a proposed altitude and wind profile for an extended flight to compare fuel burn and time with current conditions.
20	B-757	Ability to project routes that are not in use on the EFIS screen for reference purposes.
2	B-757/67	Putting winds aloft into figuring an altitude change to see how much fuel and time would be affected.
4	B-757/67	—

continued on next page...

...continued from previous page

s	Aircraft	Comments
5	B-757/67	—
6	B-757/67	I'd like to see all warning/prohibited/restricted areas displayed on the map. I'd also like to be able to display TCAS on the map.
8	B-757/67	Specific range (i.e. miles per 1000 pounds of fuel).
9	B-757/67	If FMS could forecast optimum and maximum altitudes for down the road, based on predicted fuel burn, this would enhance cruise altitude planning.
10	B-757/67	—
12	B-757/67	(1) Real time radar depiction for entire route of flight. (2) A menu of several different routes to the destination with real time winds used for a fuel computation (include pressure pattern data too). Perhaps a ride rating for each route or segments of the route also.
13	B-757/67	Continuously updated maximum altitude (we have to enter an altitude in excess of max. for it to give us the max).
21	B-757/67	Flight conditions and turbulence probabilities utilizing weather reports, winds aloft, temperature changes, and PIREPS; possibly uplinked through ACARS and indicated in a FMS format.
24	B-757/67	—
28	B-757/67	"Nothing"—system works very well and up until now I have always managed to get everything I need from the FMS in the 767/757/—.
29	B-757/67	PIREP ride report uplink, cruise altitude wind/ride trade-offs (i.e. change altitude up/down for "x" level of ride improvement if fuel burn increases less than "y").
31	B-757/67	It would be nice if it could get the winds uplinked at all times (e.g. North Atlantic). Also, with the winds it would be nice if there is some way of alerting you if there is a difference of winds that would produce turbulence/wind-shear at cruise altitude. Also in our FMS, there should be a way of showing max altitude at a given weight at all times and at what position or time instead of doing 3-4 step operation.
3	B-767	(1) GPS-based GPWS to display terrain for earlier warning. (2) Constant readout of optimum cruise altitude. (3) Easier flight planning enroute, to take advantage of changing conditions.

continued on next page...

...continued from previous page

s	Aircraft	Comments
11	B-767	Memory page for winds and ground speed at altitudes you have climbed or descended through. Page for proposed ground speed at altitude above and below based on actual winds at those altitudes.
14	B-767	Can't think of any additional features. What would help most with cruise-altitude decision making would be for ATC to open up altitudes above FL 290 to every 1000-ft. increment (i.e. FL 300, 320, 340, etc.).
18	B-767	Direct sharing of wind information between aircraft, without the intervening ground/ data-link.
19	B-767	Wind plots from forecast info, input by pilots displayed in a map overlay.
25	B-767	Ability to predict wind changes with altitude and to factor this into determining optimal altitude.
26	B-767	—
27	B-767	Additional memory capability for navigation (i.e. ability to load/display high terrain, forecast frontal systems, SIGMETs, AIRMETs information). Better ability to compute fuel burn and time with varying winds, over length of flight.
32	B-767	Real time presentation of max altitude capability next to optimum altitude. Miles per 1000 lb of fuel burn presentation.
1	MD-11	The ability to test the cost consequences of any altitude change, including new Mach, forecast wind grid (matrix) and optimal/non-optimal fuel burn.

B.2 Pilot-Dispatcher Survey

In this second survey, pilots and dispatchers were asked to rank the six attributes of the outcome most often mentioned in the first survey: ride quality, fuel burn, flight time, fuel at destination, schedule adherence, and safety. The results are shown in Tables B.3 and B.4 respectively.

<i>subj.</i>	<i>safety</i>	<i>ride quality</i>	<i>fuel at dest.</i>	<i>sched.</i>	<i>fuel burn</i>	<i>flight time</i>
<i>p1</i>	1	2	4.5	4.5	4.5	4.5
<i>p2</i>	1	2	6	4	3	5
<i>p3</i>	1	3	5	2	4	6
<i>p4</i>	1	2	6	4	3	5
<i>p5</i>	1	3	2	4	6	5
<i>p6</i>	1	4	2	3	5	6
<i>p7</i>	1	3	2	4	5	6
<i>p8</i>	1	4	2	3	5	6
<i>p9</i>	1	2	4	3	5	6
<i>p11</i>	1	2	5.5	3.5	5.5	3.5
<i>p12</i>	1	2	3	4	5	6
<i>p13</i>	1	4	2	3	5	6
<i>p14</i>	1	2	6	3	4	5
<i>p15</i>	1	3	2	4	5	6
<i>p16</i>	1	2	3	5	4	6
<i>p17</i>	1	2	3	4	5	6
<i>p18</i>	1	2	4	3	5	6
<i>p19</i>	1	3	2	5	4	6
<i>p20</i>	1	2	4	3	6	5
<i>p22</i>	1	3	2	4	5	6
<i>p23</i>	1	4	4	4	4	4
<i>p24</i>	1	2	4	5	3	6
<i>p25</i>	1	2	4	5	3	6
<i>p28</i>	1	2	4.5	3	4.5	6
<i>p29</i>	1	2	4	6	3	5
<i>p30</i>	1	2	3	6	4	5
<i>p31</i>	1	2	5	3	4	6
<i>p32</i>	1	2	4	3	6	5
<i>p33</i>	1	2	3	4	5	6
<i>p34</i>	1	2	3	5	4	6
<i>p35</i>	1	4	2	6	3	5
<i>p36</i>	1	4	3	5	2	6
<i>p37</i>	1	2	3	6	4	5
<i>p38</i>	1	3	2	5	4	6
<i>p39</i>	1	3	2	6	4	5
<i>means</i>	1	2.57	3.44	4.14	4.33	5.51

Table B.3: The rankings of six attributes of the outcome given by the subject pilots. Concordance, $W = 0.701$.

<i>subj.</i>	<i>safety</i>	<i>ride quality</i>	<i>sched.</i>	<i>fuel burn</i>	<i>fuel at dest.</i>	<i>flight time</i>
<i>d1</i>	1	2	3	5	5	5
<i>d2</i>	1	2	3	6	5	4
<i>d3</i>	1	2	3	5	6	4
<i>d4</i>	1	5	2	3	4	6
<i>d5</i>	1	2	4	6	3	5
<i>d6</i>	1	2	3	4	6	5
<i>d7</i>	1	3	2	4	5	6
<i>means</i>	1	2.57	2.86	4.71	4.86	5

Table B.4: The rankings of six attributes of the outcome given by the subject dispatchers. Concordance, $W = 0.748$.

<i>subj.</i>	PIREP	FMS	<i>flightplan</i>	POH	<i>dispatch</i>
<i>p1</i>	1.5	1.5	3	5	4
<i>p2</i>	1	2	4	3	5
<i>p3</i>	3	2	4	5	1
<i>p4</i>	1	4	2	5	3
<i>p5</i>	2	5	3	1	4
<i>p6</i>	2	4	3	1	5
<i>p7</i>	1	2	3	4	5
<i>p8</i>	1	2	3	4	5
<i>p9</i>	1	2	3	5	4
<i>p11</i>	3	1	2	5	4
<i>p12</i>	2	3	1	5	4
<i>p13</i>	2	3	5	1	4
<i>p14</i>	3	1.5	4	1.5	5
<i>p15</i>	1	2	4	3	5
<i>p16</i>	2	4	5	1	3
<i>p17</i>	1	3	4	2	5
<i>p18</i>	3	2	1	4	5
<i>p19</i>	1	4	2	5	3
<i>p20</i>	1	2	3	5	4
<i>p22</i>	3	4	2	1	5
<i>p23</i>	1	4.5	3	4.5	2
<i>p24</i>	1	4	3	5	2
<i>p25</i>	1	2	4	3	5
<i>p28</i>	1	4	3	2	5
<i>p29</i>	1	4	2	5	3
<i>p30</i>	1	3	5	2	4
<i>p31</i>	1	2	3	4	5
<i>p32</i>	1	2	3	5	4
<i>p33</i>	1	3	4	5	2
<i>p34</i>	1	4	3	2	5
<i>means</i>	1.52	2.88	3.13	3.47	4.00

Table B.5: The rankings of six sources of information given by the subject pilots. Concordance, $W = 0.345$.

Appendix C

Utility Assessment Experiment Data

This Appendix contains supplemental results from the utility assessment experiment of Chapter 4.

C.1 Conditional Utilities

Conditional utilities for each of the subjects are shown in Figures C-1 through C-5.




Figure C-1: Conditional utilities for subject one (B-747-400).




Figure C-2: Conditional utilities for subject two (MD-11).




Figure C-3: Conditional utilities for subject three (B-767).




Figure C-4: Conditional utilities for subject four (B-747-400).




Figure C-5: Conditional utilities for subject five (B-747-400).

C.2 Multilinear Utilities

C.2.1 Multilinear Utility Functions

Equation C.1 shows the multilinear utility function used to model the subjects' decisions:

$$\begin{aligned}\mathcal{U}_{ftq}(f, t, q) = & k_f \mathcal{U}_f(f) + k_t \mathcal{U}_t(t) + k_q \mathcal{U}_q(q) \\ & + k_{ft} k_f k_t \mathcal{U}_f(f) \mathcal{U}_t(t) + k_{fq} k_f k_q \mathcal{U}_f(f) \mathcal{U}_q(q) + k_{tq} k_t k_q \mathcal{U}_t(t) \mathcal{U}_q(q) \\ & + k_{ftq} k_f k_t k_q \mathcal{U}_f(f) \mathcal{U}_t(t) \mathcal{U}_q(q).\end{aligned}\tag{C.1}$$

Values of the aggregated coefficients $k_t \dots k_{ftq} k_f k_t k_q$ for each subject are shown in Table C.1.

	S_1	S_2	S_3	S_4	S_5
k_f	0.150	0.359	0.175	0.821	0.500
k_t	0.307	0.096	0.250	0.010	0.350
k_q	0.614	0.012	0.650	0.154	0.400
$k_{ft} k_f k_t$	-0.024	0.228	-0.125	0.015	-0.180
$k_{fq} k_f k_q$	0.008	0.042	-0.150	0.010	-0.275
$k_{tq} k_t k_q$	-0.008	0.036	-0.200	0.010	-0.175
$k_{ftq} k_f k_t k_q$	-0.047	0.228	0.400	-0.021	0.380

Table C.1: Multilinear utility coefficients by subject.

C.2.2 Graphical Comparison of Utility and Cost Models

Figures C-6 through C-10 show two objective functions over destination fuel and flight time: utility (with turbulence constant at zero) and cost (computed using Equation 2.2). Note that for every pilot, there is a substantial difference between the shapes of the utility and cost models, and therefore between optimal policy adjustments made according to each.




Figure C-6: A comparison of utility (with turbulence held constant at zero) and cost models for subject one (B-747-400).

◇ ◇ ◇




Figure C-7: A comparison of utility (with turbulence held constant at zero) and cost models for subject two (MD-11).




Figure C-8: A comparison of utility (with turbulence held constant at zero) and cost models for subject three (B-767).




Figure C-9: A comparison of utility (with turbulence held constant at zero) and cost models for subject four (B-747-400).




Figure C-10: A comparison of utility (with turbulence held constant at zero) and cost models for subject five (B-747-400).

Appendix D

Navigation Algorithms

The algorithms described in this chapter permit the calculation of distance and heading between waypoints, and thus form a basis for the distance and heading calculations used in the performance and optimization code. In addition, several sources of error are uncovered, and the magnitudes of the errors they generate are estimated. These errors are important to consider when defining requirements for an optimization system for the trajectory planning problem.

D.1 Coordinate Systems

Consider waypoints 1 and 2 in Figure D-1, which have latitudes ψ_1 and ψ_2 (measured in degrees North of the Equator) and longitudes ϕ_1 and ϕ_2 (measured in degrees East of the prime meridian, which runs through Greenwich, England) respectively. A cartesian coordinate system is defined with its origin at the center of the earth, its x -axis (unit vector \mathbf{i}) pointing through the equator at the prime meridian ($\psi = 0^\circ$, $\phi = 0^\circ$), its y -axis (unit vector \mathbf{j}) pointing out through the equator at the 90° East meridian ($\psi = 0^\circ$, $\phi = 90^\circ$), and its z -axis (unit vector \mathbf{k}) pointing up through the North Pole ($\psi = 90^\circ$).




Figure D-1: The Earth-centered coordinate systems, showing great-circle distance d , and initial true course ω , from waypoint 1 to waypoint 2.

D.2 Great-Circle Distance

D.2.1 Spherical-Earth Calculation

In order to calculate the great-circle distance between two waypoints, the latitude and longitude of each waypoint are first transformed into a unit vector which points to that waypoint in the aforementioned cartesian coordinate system:

$$\mathbf{u}_1 = \cos \psi_1 \cos \phi_1 \cdot \mathbf{i} + \cos \psi_1 \sin \phi_1 \cdot \mathbf{j} + \sin \psi_1 \cdot \mathbf{k}, \quad (\text{D.1})$$

$$\mathbf{u}_2 = \cos \psi_2 \cos \phi_2 \cdot \mathbf{i} + \cos \psi_2 \sin \phi_2 \cdot \mathbf{j} + \sin \psi_2 \cdot \mathbf{k}. \quad (\text{D.2})$$

The dot-product of two vectors is proportional to the cosine of the angle, γ , between them:

$$\mathbf{u}_1 \cdot \mathbf{u}_2 \equiv \|\mathbf{u}_1\| \|\mathbf{u}_2\| \cos(\gamma), \quad (\text{D.3})$$

or, since \mathbf{u}_1 and \mathbf{u}_2 are unit vectors:

$$\gamma \equiv \arccos \left[\frac{\mathbf{u}_1 \cdot \mathbf{u}_2}{\|\mathbf{u}_1\| \|\mathbf{u}_2\|} \right] = \arccos(\mathbf{u}_1 \cdot \mathbf{u}_2). \quad (\text{D.4})$$

Note that γ is never more than 180° , and thus—appropriately—defines the shorter of the two great-circle paths from waypoint 1 to waypoint 2. Assuming that the earth is a perfect sphere with radius R_e , the great-circle distance, d , measured along the surface of the earth¹ is simply:

$$d = R_e \gamma \frac{\pi}{180}. \quad (\text{D.5})$$

D.2.2 Ellipsoidal Error

In fact, the earth is not a perfect sphere. It is more closely approximated by an ellipsoid of revolution with a mean equatorial radius of 6,378 km, and a polar radius of 6,357 km (Smithsonian, 1996). The ratio of these two numbers provides an upper bound on the error introduced by making the assumption that the earth is a perfect sphere instead of an ellipsoid: it is the ratio of two short paths, one at the equator and the other at the pole, which subtend the same angle at the center of the earth, and is about 1.0033. Since neither of these radii is more correct than the other, we can assume that the best value lies midway between them, and then the maximum error due to the spherical-earth assumption would be reduced by a factor of about two, to about 0.17%.

D.2.3 Altitude-related Error

An aircraft flying at an altitude of several miles above the Earth's surface will follow a longer path than it would flying just above the Earth's surface, as shown in Figure D-2. However, many navigation devices (from hand-held GPS units to the FMSS in some large aircraft) use distance measured along the surface of the earth instead of this flight distance for their performance calculations. In order to examine the size of the resulting error, we can differentiate Equation D.5 with respect to R_e :

$$\frac{\partial d}{\partial R_e} = \gamma \frac{\pi}{180} = \frac{d}{R_e}. \quad (\text{D.6})$$

¹We use $R_e = 3437.75$ nautical miles




Figure D-2: Altitude-related error.

Then the percentage change in flight-path length is equal to the percentage change in R_e :

$$\frac{\Delta d}{d} = \frac{\Delta R_e}{R_e}. \quad (\text{D.7})$$

For an aircraft travelling at an altitude of 36,000 feet ($\Delta R_e \approx 6$ nautical miles) $\Delta d/d \approx 0.2\%$. Put another way, each additional ICAO step climb (of four thousand feet) adds approximately 0.02% to the distance which must be flown. Although small, this represents a significant error because the fuel savings from optimization are on the order of only 0.5% (see Chapter 5), and a consistent error, because the omission of an altitude correction consistently penalizes lower-altitude trajectories.

D.3 Great-Circle Course

To calculate the initial² true course for flight from waypoint 1 to waypoint 2, \mathbf{u}_2 and the unit vector which points to the geographic (or true) North Pole ($\mathbf{u}_N = \mathbf{k}$) are projected onto the plane whose normal vector is \mathbf{u}_1 :

$$\mathbf{u}'_2 = \mathbf{u}_2 - \mathbf{u}_1(\mathbf{u}_2 \cdot \mathbf{u}_1), \quad (\text{D.8})$$

$$\mathbf{u}'_N = \mathbf{u}_N - \mathbf{u}_1(\mathbf{u}_N \cdot \mathbf{u}_1). \quad (\text{D.9})$$

The angle between these projected vectors, δ , is then calculated using their dot-product in the same manner as in Equation D.4:

$$\delta \equiv \arccos \left[\frac{\mathbf{u}'_2 \cdot \mathbf{u}'_N}{\|\mathbf{u}'_2\| \|\mathbf{u}'_N\|} \right]. \quad (\text{D.10})$$

However, since the inverse cosine function maps the range $[-1, 1]$ into the range $[0^\circ, 180^\circ]$ by providing the inside angle, rather than the angle measure clockwise from North, and true course can be anywhere in the range $[0^\circ, 360^\circ]$, there remains an ambiguity. To determine whether the course is in the Eastern half or the Western half of the projection plane, we examine the sign of the Eastbound component of vector \mathbf{u}'_2 by taking its dot-product with a convenient East-positive reference vector in the projection plane, $\mathbf{u}_N \times \mathbf{u}_1$:

$$c_{east} = \mathbf{u}'_2 \cdot (\mathbf{u}_N \times \mathbf{u}_1). \quad (\text{D.11})$$

²Note that the initial and final true courses of a great circle route are rarely the same. This can be seen in Figure D-1.

If c_{east} is positive, then u'_2 is Eastbound, and $\omega = \delta$, but if c_{east} is negative, then u'_2 is Westbound, and $\omega = 360^\circ - \delta$.

$$\omega = \begin{cases} \delta & \text{if } c_{east} \geq 0, \\ 360^\circ - \delta & \text{if } c_{east} < 0. \end{cases} \quad (D.12)$$

D.4 Heading and groundspeed in the presence of wind

Once the true course w has been determined, and given a predetermined value of true airspeed, TAS , wind speed and direction must be used to find the heading to be flown, h , and the resultant ground speed, GS , as shown in Figure D-3.




Figure D-3: The wind triangle, showing the relationship between airspeed, windspeed, and groundspeed. Note that all angles are measured *clockwise* from *True North*, and that the wind angle, c , is the direction the wind comes *from*.

D.4.1 True heading

First, the inside angle, b , must be calculated from the wind angle, c , and desired true course, ω :

$$b = (180 + \omega) - c. \quad (D.13)$$

Using the sine rule, the wind correction angle,³ a , can then be calculated:

$$\frac{\sin(a)}{WS} = \frac{\sin(b)}{TAS}, \quad (D.14)$$

and

$$a = \arcsin \left[\frac{WS \cdot \sin(b)}{TAS} \right]. \quad (D.15)$$

True heading, h , is the sum of the true course and this wind correction angle:

$$h \equiv \omega + a. \quad (D.16)$$

³The wind correction angle is the angular correction which the pilot must add to the desired course to obtain the heading to be flown. Positive wind correction angles are to the right, i.e. heading is greater than course.

D.4.2 Groundspeed

Of primary importance for use in the point-mass simulation of an aircraft is groundspeed, GS . Again using the sine rule, with the triangle's third—unmarked—inside angle:

$$\frac{\sin(180 - a - b)}{GS} = \frac{\sin(b)}{TAS}. \quad (\text{D.17})$$

Rearranging, and including the pathological cases of $b = 0^\circ, 180^\circ$ for which $\sin(b) = 0$, gives:

$$GS = \begin{cases} TAS + WS & \text{if } b = 0^\circ, \\ TAS \cdot \frac{\sin(180-a-b)}{\sin(b)} & \text{if } 0^\circ < b < 180^\circ, \\ TAS - WS & \text{if } b = 180^\circ. \end{cases} \quad (\text{D.18})$$

D.4.3 Effective headwind

Finally, airlines often provide pilots with the winds aloft forecast in a simpler format than the heading and speed pairs used above. They provide a single number: effective headwind, EHW , which is defined as the difference between the scalar values of true airspeed and groundspeed:

$$EHW \equiv TAS - GS. \quad (\text{D.19})$$

In this way, pilots may see the more important effect of a wind—its effect on their groundspeed—without first having to compute the wind correction angle required to keep them on course.

◇ ◇ ◇

Appendix E

Aircraft Performance Simulation

In order to evaluate candidate trajectories, it was necessary to build an aircraft performance simulator to determine fuel burn, flight time, etc. Because actual aircraft performance data (such as drag polars, engine thrust curves, etc.) are proprietary, and closely guarded by the manufacturers, it was not possible to use them in the performance model. Instead, tabulated data from the Boeing 747-400 Performance Manual (Boeing, 1988) were used to construct a performance model of the aircraft.

Performance data for the three major phases of flight (climb, cruise, and descent) were modelled separately for use in the simulation. This section gives a brief overview of the way in which these data were used, and of their limitations.

E.1 Altitude Capability

The sequence of altitudes specified in a profile was treated as a sequence of *command* or “desired” altitudes. At each point in the simulation for which the command altitude was higher than the current altitude, maximum altitude was determined based on current weight. All climbs were made to the lower of the aircraft’s maximum altitude and the desired altitude. In this way, it was possible for the optimization algorithm to generate candidate trajectories without regard to the aircraft’s altitude capability, saving execution time and reducing program complexity.

Initially, altitude capability was modelled with a cubic least-squares fit to data from the performance manual. These maximum altitude polynomials are depicted in gray in Figure E-1. However, the maximum altitude routine was used so frequently in the simulation that it limited the execution speed of the optimization functions. Also, the simulation needed to be constrained to “fly” at only those altitudes for which there were fuel flow and airspeed data in the cruise performance tables. Implementing a check for this criterion in the table look-up functions reduced their speeds excessively. Both problems were fixed by the building of a maximum altitude function which consisted of a binary lookup tree containing the maximum altitudes by aircraft weight from the cruise tables. A stepped function representing the values of maximum altitude from the cruise tables is superimposed in black in Figure E-1.

E.2 Climb

Figures E-2 through E-4 show the time, fuel, and distance required to climb as functions of gross weight and target altitude. Fuel and distance data were used as they appear in these figures, but the time data had been discretized (by up to 10%) when it was rounded to the nearest minute for




Figure E-1: Altitude capability as a function of weight. The solid and the continuous lines represent cubic least-squares lines fitted to the data from the performance manual for two different temperatures, while the stepped function represents the maximum altitude used in the simulation.

printing in the performance tables. To remove this discretization, a seventh-degree polynomial in altitude was used to model the relationship between climb time and altitude at each of the 21 takeoff weights (see the example for 560 klb in Figure E-5). Note that although the polynomial “misbehaves” between the published values, it fits the points themselves well, and removes the discretizations. The 21 polynomials thus obtained were then used to reconstruct values of climb time at the published altitude points. The lower panel of Figure E-2 shows the results of this smoothing process. The reconstructed values of climb time and the unaltered values of climb fuel and climb distance were placed in look-up tables for use in the simulation. The climb was simulated in 1000-foot intervals, with groundspeed calculated in the manner outlined in Section D.4 based on the prevailing winds at each thousand foot level. An estimate of the climb velocity was obtained by dividing the distance for each 1000-foot climb by the time taken.

E.3 Cruise

To simulate cruising flight, the aircraft’s altitude, A , was held constant. Fuel flow, \dot{f} , and true airspeed, TAS , were obtained from lookup tables, which were based on the data shown in the upper and lower panels respectively of Figure E-6. True airspeed and wind speed were then combined in the manner outlined in Section D.4 to produce wind correction angle (which was incidental to this work) and ground speed, GS . A forward difference integration scheme was used to update the vehicle mass, m , and distance flown, D , from time t_k to time t_{k+1} :

$$m_{k+1} = m_k - \dot{f} \cdot \Delta t, \quad (\text{E.1})$$

$$D_{k+1} = D_k + GS \cdot \Delta t, \quad (\text{E.2})$$

where $\Delta t = t_{k+1} - t_k$.

E.4 Descent

Fuel, time, and distance for the descent were computed from the tabular descent data in the Operations Manual. The descent time values were particularly noisy, since each entry had been rounded to the nearest minute (this effect is visible in the step-like series of dots in the top panel of Figure E-7). The data were smoothed using quartic polynomials of altitude in thousands of feet, A , whose coefficients $\alpha_0 \dots \gamma_4$ were found using an ordinary least-squares estimation. Terms which were not statistically significantly different from zero were dropped (only α_4 and γ_2 were this close to zero), leaving the following polynomials in A :

$$\hat{t}_d = \alpha_0 + \alpha_1 A + \alpha_2 A^2 + \alpha_3 A^3 \quad (\text{E.3})$$

$$\hat{f}_d = \beta_0 + \beta_1 A + \beta_2 A^2 + \beta_3 A^3 + \beta_4 A^4 \quad (\text{E.4})$$

$$\hat{d}_d = \gamma_0 + \gamma_1 A + \gamma_3 A^3 + \gamma_4 A^4 \quad (\text{E.5})$$

The regressions were then rerun to give the coefficients shown in Table E.1. Figure E-7 shows the

coeff.	time	fuel	distance
A^0	$+2.557 \times 10^0$	$+7.300 \times 10^{-1}$	$+5.818 \times 10^0$
A^1	$+9.630 \times 10^{-1}$	$+9.391 \times 10^{-2}$	$+3.712 \times 10^0$
A^2	-2.050×10^{-2}	-3.483×10^{-3}	0*
A^3	$+2.033 \times 10^{-4}$	$+6.371 \times 10^{-5}$	-9.654×10^{-4}
A^4	0*	-4.068×10^{-7}	$+1.232 \times 10^{-5}$
R^2	0.998	0.999	0.999

Table E.1: Coefficients of the powers of altitude, A , in the models for descent time, fuel, and distance. Also shown are the coefficients of determination, R^2 , for each model. *Starred coefficients were constrained to be zero.

performance manual data, and the polynomial models used in the simulation. Given the models of Equations E.5 and E.3, true airspeed¹ in the descent, V_d , was determined thus:

$$V_d \approx \frac{d\hat{d}_d/dA}{d\hat{t}_d/dA} = \frac{\gamma_1 + 3\gamma_3 A^2 + 4\gamma_4 A^3}{\alpha_1 + 2\alpha_2 A + 3\alpha_3 A^2} \quad (\text{E.6})$$

Groundspeed during the descent was calculated in the same manner as for the climb, combining V_d with winds at each thousand-foot interval.

E.5 Implementation

E.5.1 Programming

The performance model was prototyped using MATLAB. An object-oriented implementation was built in ANSI-standard C++ (Stroustrup, 1991) using Symantec's development system (Symantec, 1995).

¹We ignore the vertical component of the aircraft's speed in this approximation. For a typical 3° glidepath, this error only amounts to about 0.14%.

E.5.2 Choice of Integration Interval

Figure E-9 shows the percentage overestimation of fuel consumption against cruise integration interval for an entire flight. Based on these data, an interval of 0.1 hours was chosen as a compromise between accuracy and speed for use during optimization. Although it resulted in a 0.5% to 1% overestimation of fuel consumption, and a similar error in calculated flight time, it allowed the GA-based optimization to be run in a reasonable amount of time.

◇ ◇ ◇





Figure E-2: Time to climb and adjusted time to climb (in minutes) as functions of takeoff weight and target altitude. Note the ragged nature of the surface in the upper panel, which results from table entries being rounded to the nearest minute, and that in the lower panel most of this discretization has been removed.




Figure E-3: Fuel to climb (in klb) as a function of takeoff weight and target altitude.




Figure E-4: Distance to climb (in nautical miles) as a function of takeoff weight and target altitude.




Figure E-5: Time to climb as a function of target altitude for a takeoff weight of 560 klb. The dots represent values from the performance manual, while the gray line represents the seventh-order least-squares fit used to remove the discretization from the tabulated numbers. Note that this fitted line was only used to recreate values of climb time at the tabulated values of target altitude: the “hump” at about 4 kft does not interfere with the model.





Figure E-6: Cruise fuel flow (in thousands of pounds per hour) and true airspeed (in knots) as functions of weight and altitude.




Figure E-7: Time, fuel, and distance to descend, each as a function of initial altitude. Speeds: M0.88/340kt/250kt.




Figure E-8: True airspeed during the descent, as a function of altitude. The descent is nominally flown at Mach 0.88, then 340 kt, and finally 250 kt below 10,000 ft. The fine lines are hypothetical curves for Mach 0.88, 340 kt, and 250 kt. The bold line represents the speeds obtained from Equation E.6.




Figure E-9: Percentage error in estimated fuel burn, as a function of cruise integration time-interval. Note that the fuel burn for the 3.2×10^{-6} hour interval was used as the baseline for the calculation of percentage error.

References

- Asimov, I. (1970). *I, Robot*. Fawcett Publications, Inc., Greenwich, CT.
- Askey, S. Y. (1995). *Design and Evaluation of Decision Aids for Control of High-Speed Trains: Experiments and Model*. Ph.D. thesis, Department of Mechanical Engineering, M.I.T., Cambridge, MA.
- Barrows, A. K. (1993). *Development and Inflight Validation of an Automated Flight Planning System*. . . . Master's thesis, Dept. of Aeronautics and Astronautics, M.I.T., Cambridge, MA.
- Bassett, E. W. (1995). IBM's IBM Fix. *Industrial Computing*, 24-26.
- Beatty, R. (1995). *Airline Operational Control Overview*. Draft report 95133-1, Airline Dispatchers Federation, Washington, D.C.; and Seagull Technology, Inc., Cupertino, CA.
- Bellman, R. E., & Dreyfus, S. E. (1962). *Applied Dynamic Programming*. Princeton University Press, Princeton, NJ.
- Bertsekas, D. P. (1995). *Dynamic Programming and Optimal Control*, Vol. 1 of *The Art of Computer Programming*. Athena Scientific, Belmont, MA.
- Boeing (1988). *Boeing 747-400 Operations Manual, Vol 1-3*. Boeing Co., Seattle, WA.
- Chandra, D. (1989). *An Evaluation of Automation for Flight Path Management in Transport Category Aircraft*. Master's thesis, Dept. of Aeronautics and Astronautics, M.I.T., Cambridge, MA.
- Corker, K. M., & Smith, B. R. (1993). An Architecture and Model for Cognitive Engineering, Simulation and Analysis: Application to Advanced Aviation Automation. In *the proceedings of the Ninth AIAA Computing in Aerospace Conference*, pp. 1079-1088. San Diego, CA.
- Crichton, M. (1988). *Travels*. Ballantine, New York, NY.
- de Neufville, R. (1990). *Applied Systems Analysis: Engineering Planning and Technology Management*. McGraw-Hill, New York, NY.
- Delahaye, D., Alliot, J.-M., Schroenauer, M., & Farges, J.-L. (1994). Genetic Algorithms for Air Traffic Assignment. In *Proceedings of 11th European Conference on Artificial Intelligence*.
- Durand, N., Alliot, J.-M., & Chansou, O. (1995). Optimal Resolution of En Route Conflicts. *Air Traffic Control Quarterly*, 3(3), 139-161.
- Dyer, J. S. (1990). Remarks on the Analytic Hierarchy Process. *Management Science*, 36(3), 249-258.
- Federal Aviation Administration (1996a). *Aeronautical Information Manual*. Jeppeson Sanderson, Inc., Englewood, CO.
- Federal Aviation Administration (1996b). *Federal Aviation Regulations, Parts 91 and 121*. Jeppeson Sanderson, Inc., Englewood, CO.

- Gifford, R. B. (1996). Personal Communication. Unpublished.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, MA.
- Goosens, M., Mittlebach, F., & Samarin, A. (1994). *The L^AT_EX Companion*. Addison-Wesley, Reading, MA.
- Grandeau, S. C. (1995). *The Process of Airline Operational Control*. Master's thesis, Dept. of Aeronautics and Astronautics, M.I.T., Cambridge, MA.
- Honeywell (1994). *Boeing 747-400 Flight Management System Pilot's Guide*. Honeywell Inc., Phoenix, AZ.
- Jeppeson (1996a). *Instrument Rating Manual*. Jeppeson Sanderson, Inc., Englewood, CO.
- Jeppeson (1996b). *Private Pilot Manual*. Jeppeson Sanderson, Inc., Englewood, CO.
- Katz, A. (1994). *Subsonic Airplane Performance*. Society of Automotive Engineers, Inc, Warrendale, PA.
- Keeney, R. L. (1968). Evaluating Multidimensional Situations Using a Quasi-Separable Utility Function. *IEEE Transactions on Man-Machine Systems*, MMS-9(2).
- Keeney, R. L. (1995). Understanding life-threatening risks. *Risk Analysis*, 15(6), 627.
- Keeney, R. L., & Raiffa, H. (1993). *Decisions with Multiple Objectives*. Cambridge University Press, Cambridge, UK.
- Klein, G. (1993). *Naturalistic Decision Making: Implications for Design*. CSERIAC SOAR 93-01, Defense Logistics Agency, Cameron Station, Alexandria, VA.
- Lamport, L. (1994). *L^AT_EX: A Document Preparation System* (Second edition). Addison-Wesley, Reading, MA.
- Lee, H. Q., & Erzberger, H. (1980). *Algorithm for Fixed-Range Optimal Trajectories*. Technical report NASA-TP-1565, A-8003, National Aeronautics and Space Administration, Ames Research Center, Moffett Field, CA.
- Lidén, S. (1992). Optimum Cruise Profiles in the Presence of Winds. In *Proceedings of the IEEE/AIAA 11th Digital Aerospace Systems Conference*, pp. 254–261. Seattle, WA.
- Midkiff, A. H. (1996). Personal Communication. Unpublished.
- Niiya, C. K. (1990). *An Application of the A* Search Technique to Trajectory Optimization*. Master's thesis, Dept. of Aeronautics and Astronautics, M.I.T., Cambridge, MA.
- Nolan-Proxmire, D., Spracher, M., Hopkins, J., & Nelson, C. (1996). *New Cockpit Weather System to Enhance Safety and Efficiency of Airline Transport*. Press release, National Aeronautics and Space Administration, Washington, DC.
- Orasanu, J., Wich, M., Fischer, U., & Jobe, K. (1993). Distributed Problem Solving by Pilots and Dispatchers. In *the proceedings of the Seventh International Symposium on Aviation Psychology*, pp. 198–203. Columbus, OH.
- Patrick, N. J. M. (1993). *Modelling Decision-Making by Pilots*. Progress report NASA-CR-194767, Human-Machine Systems Laboratory, Massachusetts Institute of Technology, Cambridge, MA.
- Patrick, N. J. M. (1995). *Flight Trajectory Selection: Decision Theory and Optimization*. Doctoral Thesis Proposal, Department of Mechanical Engineering, M.I.T., Cambridge, MA.

- Patrick, N. J. M. (1996). *Poisson Regression Analysis of Fatal International Air-Carrier Accidents from 1976-1986*. Unpublished.
- Piper Aircraft Corporation (1978). *Arrow IV Pilot's Operating Handbook*. Piper Aircraft Corp., Vero Beach, FL.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1988). *Numerical Recipes in C*. Cambridge University Press, New York, NY.
- Roseborough, J. B. (1988). *Aiding Human Operators with State Estimates*. Ph.D. thesis, Department of Mechanical Engineering, M.I.T., Cambridge, MA.
- RTCA (1994). *Report of the RTCA Board of Directors' Select Committee on Free Flight*, 16 Dec, 1994. Tech. rep., Radio Technical Corp. of America.
- Rubbert, P. E. (1994). C.F.D. and the Changing World of Airplane Design. In *Proceedings of the 19th Congress of the International Council of Aeronautical Sciences*. Anaheim, CA.
- Rudolph, F. M., Homoki, D. A., & Sexton, G. A. (1990). *Diverter: Decision Aiding For In-Flight Diversions*. NASA contractor report 182070, Lockheed Aeronautical Systems Co., Marietta, GA.
- Saaty, T. L. (1977). A Scaling Method for Priorities in Hierarchical Structures. *Journal of Mathematical Psychology*, 15, 234-281.
- Saaty, T. L. (1990). How to Make a Decision: the Analytic Hierarchy Process. *European Journal of Operational Research*, 48(1), 9-26.
- Schreur, J. (1995). B-737 Flight Management Computer Flight Plan Trajectory Computation and Analysis. In *Proceedings of the 1995 American Controls Conference*. Seattle, WA.
- Schultz, A. C. (1991). Using a Genetic Algorithm to Learn Strategies for Collision Avoidance and Local Navigation. In *Proceedings of the Seventh International Symposium on Unmanned Untethered Submersible Technology*, pp. 213+ Marine Systems Engineering Laboratory, University of New Hampshire.
- Seigel, S., & Castellan, Jr., N. J. (1988). *Nonparametric Statistics for the Behavioral Sciences*. McGraw-Hill, New York, NY.
- Shaoee, H., & Bryson, A. E. (1976). *Airplane Minimum Fuel Flight Paths for Fixed Range*. Master's thesis, Stanford University, Stanford, CA.
- Smith, P., McCoy, E., & Orasanu, J. (1994). *Cooperative Problem Solving in the Interactions of Airline Dispatchers with ATCSCC*. Annual report NCC 2-827, NASA Ames Research Center, Moffett Field, CA.
- Smithsonian (1996). *Earth Fact Sheet*. Center for Earth and Planetary Studies, National Air and Space Museum, Smithsonian Institution, <http://ceps.nasm.edu:2020/RPIF/EARTH/earthfacts.html>.
- Stroustrup, B. (1991). *The C++ Programming Language* (Second edition). Addison-Wesley, Reading, MA.
- Sykes, J. B. (Ed.). (1982). *The Concise Oxford Dictionary of Current English* (Seventh edition). The Clarendon Press, Oxford, England.
- Symantec (1995). *Symantec C++ Development System for Power Macintosh, Version 8*. Symantec Corporation, Cupertino, CA.
- Thomas, C. L. (Ed.). (1993). *Taber's Cyclopedic Medical Dictionary* (Seventeenth edition). F.A. Davis Company, Philadelphia, PA.

- Trujillo, A. C. (1996). *Uncertainties that Flight Crews and Dispatchers Must Consider When Calculating the Fuel Needed for a Flight*. Technical memorandum NASA TM-110240, National Aeronautics and Space Administration, Langley Research Center, Hampton, VA.
- U.S. Geological Survey (1996). *Coastline Extractor*. <http://crusty.er.usgs.gov/coast/getcoast.html>.
- van Deventer, P. (1993). *Real-Time Trajectory Optimization Using a Constrained Genetic Algorithm*. Master's thesis, Department of Aeronautics and Astronautics, M.I.T., Cambridge, MA.
- Wagenmakers, J. (1991). *Aircraft Performance Engineering*. Prentice Hall, New York, NY.
- Wiener, E., Chidester, T., Kanki, B., Palmer, E., Curry, R., & Gregorich, S. (1991). *The Impact of Cockpit Automation on Crew Coordination: 1. Overview etc. . . .* Report 177587, NASA Ames Research Center, Mountain View, CA.
- Winston, P. H. (1992). *Artificial Intelligence*. Addison-Wesley, Reading, MA.
- Yang, L. C., & Hansman, Jr., R. J. (1995). Application of the Analytic Hierarchy Process for Making Subjective Comparisons Between Multiple Automation/Display Options. In *The 6th IFAC Symposium on Analysis, Design, and Evaluation of Man-Machine Systems*, pp. 555-559 Cambridge, MA.
- Yntema, D. B., & Klem, L. (1965). Telling a Computer How to Evaluate Multidimensional Situations. *IEEE Transactions on Human Factors in Electronics*, HFE-6(1).

